



MetroCluster Documentation

ONTAP MetroCluster

NetApp
February 13, 2026

This PDF was generated from <https://docs.netapp.com/us-en/ontap-metrocluster/index.html> on February 13, 2026. Always check docs.netapp.com for the latest.

Table of Contents

MetroCluster Documentation	1
MetroCluster release notes	2
What's new in MetroCluster features	2
What's new in MetroCluster IP platform and switch support	7
Platform support	7
Switch support	8
What's new in MetroCluster FC platform and switch support	8
Platform support	8
Switch support	9
What's new in ONTAP Mediator support for MetroCluster IP	9
What's new in MetroCluster Tiebreaker support	10
Enhancements	10
OS support matrix	10
Install a fabric-attached MetroCluster	12
Overview	12
Prepare for the MetroCluster installation	12
Differences among the ONTAP MetroCluster configurations	12
Cluster peering	13
Considerations for ISLs	16
Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations	17
Requirements for using a Brocade DCX 8510-8 switch	18
Considerations when using unmirrored aggregates	19
Firewall usage at MetroCluster sites	20
Cable a fabric-attached MetroCluster configuration	21
Cabling a fabric-attached MetroCluster configuration	21
Parts of a fabric MetroCluster configuration	21
Required MetroCluster FC components and naming conventions	28
Configuration worksheets for FC switches and FC-to-SAS bridges	31
Install and cable MetroCluster components	31
Configure the FC switches	67
Install FC-to-SAS bridges and SAS disk shelves	192
Configure the MetroCluster FC software in ONTAP	206
Gathering required information	207
Similarities and differences between standard cluster and MetroCluster configurations	213
Verifying and configuring the HA state of components in Maintenance mode	214
Restoring system defaults and configuring the HBA type on a controller module	215
Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems	217
Verifying disk assignment in Maintenance mode in an eight-node or a four-node configuration	219
Verifying disk assignment in Maintenance mode in a two-node configuration	226
Setting up ONTAP	227
Configuring the clusters into a MetroCluster configuration	233
Checking for MetroCluster configuration errors with Config Advisor	263
Verifying local HA operation	264

Verifying switchover, healing, and switchback	266
Protecting configuration backup files	266
Considerations for using virtual IP and Border Gateway Protocol with a MetroCluster configuration	266
ONTAP limitations	267
Guidelines for using this Layer 3 solution with a MetroCluster configuration	268
Testing the MetroCluster configuration	268
Verifying negotiated switchover	269
Verifying healing and manual switchback	270
Loss of a single FC-to-SAS bridge	273
Verifying operation after power line disruption	275
Verifying operation after a switch fabric failure	276
Verifying operation after loss of a single storage shelf	278
Remove MetroCluster configurations	288
How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring	289
Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring	289
Synchronize the system time using NTP	289
Considerations when using ONTAP in a MetroCluster configuration	290
FlexCache support in a MetroCluster configuration	291
FabricPool support in MetroCluster configurations	291
FlexGroup support in MetroCluster configurations	292
Consistency group support in MetroCluster configurations	292
Job schedules in a MetroCluster configuration	292
Cluster peering from the MetroCluster site to a third cluster	293
LDAP client configuration replication in a MetroCluster configuration	293
Networking and LIF creation guidelines for MetroCluster configurations	293
SVM disaster recovery in a MetroCluster configuration	298
Output for the "storage aggregate plex show" command is indeterminate after a MetroCluster switchover	301
Modifying volumes to set the NVFAIL flag in case of switchover	301
Where to find additional information	302
MetroCluster and miscellaneous information	302
Install a MetroCluster IP configuration	304
MetroCluster IP installation workflow	304
Prepare for the MetroCluster installation	304
ONTAP MetroCluster configurations support matrix	304
Differences between ONTAP Mediator and MetroCluster Tiebreaker	305
Learn about remote storage and MetroCluster IP configurations	306
MetroCluster IP requirements for automatic drive assignment and ADP systems	308
Requirements for cluster peering in MetroCluster IP configurations	324
ISL requirements	326
Considerations for using MetroCluster-compliant switches	341
Learn about unmirrored aggregates in MetroCluster IP configurations	349
Firewall port requirements for MetroCluster IP configurations	350

Learn about using virtual IP and Border Gateway Protocol with a MetroCluster IP configuration	351
Configure the MetroCluster hardware components	353
Learn about hardware component interconnections in a MetroCluster IP configuration	354
Required MetroCluster IP configuration components and naming conventions	358
Rack the MetroCluster IP configuration hardware components.	362
Cable the MetroCluster IP switches	362
Cable the ONTAP controller module ports in a MetroCluster IP configuration.	414
Configure the MetroCluster IP switches.	415
Monitor MetroCluster IP switch health	472
Configure the MetroCluster software in ONTAP	499
Configure the MetroCluster software using the CLI	499
Configure the MetroCluster software using System Manager	566
Configure ONTAP Mediator for unplanned automatic switchover	569
ONTAP Mediator installation requirements for MetroCluster IP configurations	569
Set up the ONTAP Mediator for a MetroCluster IP configuration.	572
Remove the ONTAP Mediator from a MetroCluster IP configuration.	576
Connect a MetroCluster IP configuration to a different ONTAP Mediator instance	577
How the ONTAP Mediator supports automatic unplanned switchover in MetroCluster IP configurations	577
Manage the ONTAP Mediator with System Manager in MetroCluster IP configurations	579
Test the ONTAP node switchover for your MetroCluster IP configuration	580
Verifying negotiated switchover	580
Verifying healing and manual switchover	582
Verifying operation after power line disruption	585
Verifying operation after loss of a single storage shelf.	587
Remove MetroCluster configurations	597
Requirements and considerations for ONTAP operations with MetroCluster IP configurations.	598
Licensing considerations	598
SnapMirror consideration	598
MetroCluster operations in ONTAP System Manager	598
FlexCache support in a MetroCluster configuration.	598
FabricPool support in MetroCluster configurations	599
FlexGroup support in MetroCluster configurations	600
Job schedules in a MetroCluster configuration	600
Cluster peering from the MetroCluster site to a third cluster	600
LDAP client configuration replication in a MetroCluster configuration	600
Networking and LIF creation guidelines for MetroCluster configurations.	600
SVM disaster recovery in a MetroCluster configuration.	604
Output for the storage aggregate plex show command is indeterminate after a MetroCluster switchover	607
Modifying volumes to set the NVFAIL flag in case of switchover.	607
How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring.	608
Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring in a MetroCluster IP configuration	608
Synchronize the system time using NTP in a MetroCluster IP configuration	608

Where to find additional information about MetroCluster IP	609
MetroCluster and miscellaneous information	609
Install a stretch MetroCluster configuration	612
Overview	612
Prepare for the MetroCluster installation	612
Differences between the ONTAP MetroCluster configurations	612
Cluster peering	614
Considerations when using unmirrored aggregates	616
Firewall usage at MetroCluster sites	617
Choosing the correct installation procedure for your configuration	617
Cable a two-node SAS-attached stretch MetroCluster configuration	618
Cabling a two-node SAS-attached stretch MetroCluster configuration	618
Parts of a two-node SAS-attached stretch MetroCluster configuration	618
Required MetroCluster hardware components and naming guidelines for two-node SAS-attached stretch configurations	619
Install and cable MetroCluster components for two-node SAS-attached stretch configurations	620
Cable a two-node bridge-attached stretch MetroCluster configuration	623
Cabling a two-node bridge-attached stretch MetroCluster configuration	623
Parts of a two-node bridge-attached stretch MetroCluster configuration	623
Required MetroCluster hardware components and naming conventions for two-node bridge-attached stretch configurations	624
Information gathering worksheet for FC-to-SAS bridges	626
Install and cable MetroCluster components	628
Install FC-to-SAS bridges and SAS disk shelves	630
Configuring the MetroCluster software in ONTAP	643
IP network information worksheet for Site A	644
IP network information worksheet for site B	646
Similarities and differences between standard cluster and MetroCluster configurations	648
Restoring system defaults and configuring the HBA type on a controller module	649
Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems	651
Verifying disk assignment in Maintenance mode in a two-node configuration	653
Verifying the HA state of components	655
Setting up ONTAP in a two-node MetroCluster configuration	656
Configuring the clusters into a MetroCluster configuration	658
Checking for MetroCluster configuration errors with Config Advisor	681
Verifying switchover, healing, and switchback	681
Protecting configuration backup files	682
Considerations for using virtual IP and Border Gateway Protocol with a MetroCluster configuration	682
Testing the MetroCluster configuration	684
Verifying negotiated switchover	685
Verifying healing and manual switchback	686
Loss of a single FC-to-SAS bridge	689
Verifying operation after power line disruption	691
Verifying operation after loss of a single storage shelf	692
Remove MetroCluster configurations	703

How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring	704
Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring	704
Synchronize the system time using NTP	704
Considerations when using ONTAP in a MetroCluster configuration	705
Licensing considerations	706
SnapMirror consideration	706
FlexCache support in a MetroCluster configuration	706
FabricPool support in MetroCluster configurations	706
FlexGroup support in MetroCluster configurations	707
Job schedules in a MetroCluster configuration	707
Cluster peering from the MetroCluster site to a third cluster	707
LDAP client configuration replication in a MetroCluster configuration	708
Networking and LIF creation guidelines for MetroCluster configurations	708
SVM disaster recovery in a MetroCluster configuration	712
Output of the storage disk show and storage shelf show commands in a two-node stretch MetroCluster configuration	714
Output for the storage aggregate plex show command is indeterminate after a MetroCluster switchover	714
Modifying volumes to set the NVFAIL flag in case of switchover	715
Transitioning from a stretch to a fabric-attached MetroCluster configuration	715
Where to find additional information	716
MetroCluster and miscellaneous information	716
Install and configure MetroCluster Tiebreaker	719
What's new in MetroCluster Tiebreaker support	719
Enhancements	719
OS support matrix	719
Overview of the Tiebreaker software	720
Detecting failures with NetApp MetroCluster Tiebreaker software	720
How the Tiebreaker software detects site failures	721
How the Tiebreaker software detects intersite connectivity failures	721
How different disaster types affect Tiebreaker software detection time	722
About the Tiebreaker CLI and man pages	723
Install the Tiebreaker software	723
Tiebreaker installation workflow	723
Prepare to install the Tiebreaker software	724
Secure the Tiebreaker host and database installation	725
Install the Tiebreaker software package	728
Upgrade the host where the Tiebreaker monitor is running	822
Configure the MetroCluster Tiebreaker software	822
Launch the Tiebreaker software CLI	822
Add MetroCluster configurations	822
Import certificates	826
Commands for modifying MetroCluster Tiebreaker configurations	830

Remove MetroCluster configurations	831
Configuring SNMP settings for Tiebreaker software	832
Monitoring the MetroCluster configuration	834
Configuring AutoSupport	834
Displaying the status of monitoring operations	836
Displaying MetroCluster configuration information	838
Creating dump files	838
Disable Tiebreaker observer mode	839
Risks and limitations of using MetroCluster Tiebreaker in active mode	839
Firewall requirements for MetroCluster Tiebreaker	840
Simulate a switchover using MetroCluster Tiebreaker	840
Event log files for MetroCluster Tiebreaker	842
Where to find additional information	842
MetroCluster and miscellaneous information	842
Understand MetroCluster data protection and disaster recovery	844
Understanding MetroCluster data protection and disaster recovery	844
How eight- and four-node MetroCluster configurations provide local failover and switchover	844
How MetroCluster configurations provide data and configuration replication	845
Types of disasters and recovery methods	851
How an eight-node or four-node MetroCluster configuration provides nondisruptive operations	853
How a two-node MetroCluster configuration provides nondisruptive operations	854
Overview of the switchover process	854
What happens during healing (MetroCluster FC configurations)	861
What happens during healing (MetroCluster IP configurations)	862
Automatic healing of aggregates on MetroCluster IP configurations after switchover	862
Creating SVMs for a MetroCluster configuration	865
What happens during a switchback	866
Perform switchover, healing, and switchback	867
Perform switchover for tests or maintenance	867
Performing switchover for tests or maintenance	867
Limitations when the MetroCluster configuration is in switchover	867
Verifying that your system is ready for a switchover	868
Sending a custom AutoSupport message prior to negotiated switchover	869
Performing a negotiated switchover	869
Verify that the SVMs are running and the aggregates are online	871
Heal the configuration	872
Performing a switchback	876
Verifying a successful switchback	879
Commands for switchover, healing, and switchback	880
Use System Manager to perform switchover and switchback (MetroCluster IP configurations only)	881
Overview of switchover and switchback	881
Use System Manager in ONTAP 9.6 or 9.7 for switchover and switchback	881
Use System Manager in ONTAP 9.8 or later for switchover and switchback	882
Monitoring the MetroCluster configuration	883
Checking the MetroCluster configuration	883

Commands for checking and monitoring the MetroCluster configuration	886
Using the MetroCluster Tiebreaker or ONTAP Mediator to monitor the configuration	887
How the NetApp MetroCluster Tiebreaker software detects failures	887
Monitoring and protecting the file system consistency using NVFAIL	889
How NVFAIL impacts access to NFS volumes or LUNs	889
Commands for monitoring data loss events	890
Accessing volumes in NVFAIL state after a switchover	891
Recovering LUNs in NVFAIL states after switchover	891
Where to find additional information	892
MetroCluster and miscellaneous information	892
Maintain the MetroCluster components	894
Learn about MetroCluster maintenance	894
Prepare for maintenance tasks	894
Maintenance procedures for different types of MetroCluster configurations	894
All other maintenance procedures	894
Prepare for MetroCluster maintenance	895
Enable console logging before performing maintenance tasks	895
Remove ONTAP Mediator or Tiebreaker monitoring before performing maintenance tasks	896
MetroCluster failure and recovery scenarios	897
Using the Interoperability Matrix Tool to find MetroCluster information	898
Maintenance procedures for MetroCluster FC configurations	898
Modify a switch or ATTO bridge IP address for health monitoring	899
FC-to-SAS bridge maintenance	900
FC switch maintenance and replacement	959
Replacing a shelf nondisruptively in a fabric-attached MetroCluster configuration	1010
Hot add storage to a MetroCluster FC configuration	1016
Hot-removing storage from a MetroCluster FC configuration	1038
Power off and power on a single site in a MetroCluster FC configuration	1042
Powering off an entire MetroCluster FC configuration	1056
Maintenance procedures for MetroCluster IP configurations	1058
IP switch maintenance and replacement	1058
Identifying storage in a MetroCluster IP configuration	1084
Adding shelves to a MetroCluster IP using shared Storage MetroCluster switches	1088
Configure end-to-end encryption in a MetroCluster IP configuration	1104
Power off and power on a single site in a MetroCluster IP configuration	1108
Powering off an entire MetroCluster IP configuration	1115
Maintenance procedures for all MetroCluster configurations	1117
Replacing a shelf nondisruptively in a stretch MetroCluster configuration	1117
When to migrate root volumes to a new destination	1119
Moving a metadata volume in MetroCluster configurations	1119
Renaming a cluster in MetroCluster configurations	1122
Verify the health of a MetroCluster configuration	1124
Where to find additional information	1126
Transition from MetroCluster FC to MetroCluster IP	1128
Choose your transition procedure	1128

Supported platform combinations	1128
Choose your transition procedure	1130
Transition nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later)	1131
Transitioning nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later)	1131
Prepare for transition from a MetroCluster FC to a MetroCluster IP configuration	1132
Transition from MetroCluster FC to MetroCluster IP configurations	1143
Sending a custom AutoSupport message after maintenance	1197
Restoring Tiebreaker or Mediator monitoring	1197
Disruptively transition from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)	1197
Disruptively transitioning from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)	1197
Example naming in this procedure	1198
Preparing for disruptive FC-to-IP transition	1199
Transitioning the MetroCluster FC nodes	1208
Connecting the MetroCluster IP controller modules	1217
Configuring the new nodes and completing transition	1233
Returning the system to normal operation	1238
Disruptively transitioning from MetroCluster FC to MetroCluster IP when retiring storage shelves (ONTAP 9.8 and later)	1240
Enable console logging	1240
Requirements for transition when retiring old shelves	1241
Workflow for disruptive transition when moving data and retiring old storage shelves	1241
Transitioning the configuration	1242
Migrating the root aggregates	1243
Migrating the data aggregates	1244
Retiring shelves moved from node_A_1-FC and node_A_2-FC	1244
Completing transition	1245
Disruptively transitioning when existing shelves are not supported on new controllers (ONTAP 9.8 and later)	1246
Enable console logging	1246
Requirements for transition when shelves are not supported on the new nodes	1246
Workflow for disruptive transition when shelves are not supported by new controllers	1247
Preparing the new controller modules	1248
Attaching the new disk shelves to the existing MetroCluster FC controllers	1248
Migrate root aggregates and move data to the new disk shelves	1249
Transitioning the configuration	1255
Moving an FC SAN workload from MetroCluster FC to MetroCluster IP nodes	1256
Move an FC SAN workload from MetroCluster FC to MetroCluster IP nodes	1256
Move Linux iSCSI hosts from MetroCluster FC to MetroCluster IP nodes	1263
Step 1: Set up new iSCSI connections	1263
Step 2: Add the new nodes as reporting nodes	1267
Step 3: Remove reporting nodes and rescan paths	1272

Where to find additional information	1274
MetroCluster and miscellaneous information	1274
Upgrade, refresh, or expand the MetroCluster configuration	1277
Start here - Choose your procedure	1277
Start here: Choose between controller upgrade, system refresh, or expansion	1277
Choose a controller upgrade procedure	1278
Choosing a system refresh method	1284
Choose an expansion procedure	1288
Upgrade controllers in four-node MetroCluster IP using switchover and switchback with "system controller replace" commands (ONTAP 9.13.1 and later)	1289
Workflow for upgrading MetroCluster IP controllers using "system controller replace" commands (ONTAP 9.13.1 or later)	1289
Prepare to upgrade	1290
Upgrade your controllers	1302
Complete the MetroCluster IP controller upgrade	1327
Upgrade controllers in MetroCluster IP using switchover and switchback (ONTAP 9.8 and later)	1328
Workflow for MetroCluster IP controller upgrades using switchover and switchback (ONTAP 9.8 and later)	1328
Prepare to upgrade	1329
Upgrade your controllers	1343
Complete the MetroCluster IP controller upgrade	1371
Upgrade controllers from AFF A700/FAS9000 to AFF A900/FAS9500 in a MetroCluster IP configuration using switchover and switchback (ONTAP 9.10.1 or later)	1373
Enable console logging	1373
Workflow for upgrading controllers in a MetroCluster IP configuration	1374
Prepare for the upgrade	1374
Switch over the MetroCluster configuration	1380
Remove AFF A700 or FAS9000 platform controller module and NVS	1382
Install the AFF A900 or FAS9500 NVS and controller modules	1383
Update the switch RCF files to accommodate the new platforms	1390
Configure the new controllers	1390
Switch back the MetroCluster configuration	1399
Check the health of the MetroCluster configuration	1401
Upgrade the nodes on site_A	1402
Restore Tiebreaker or Mediator monitoring	1402
Send a custom AutoSupport message after maintenance	1403
Upgrade controllers in a MetroCluster FC configuration using switchover and switchback	1403
Supported platform combinations	1403
About this task	1403
Enable console logging	1404
Prepare for the upgrade	1405
Switch over the MetroCluster configuration	1411
Prepare the network configuration of the old controllers	1413
Remove the old platforms	1415
Configure the new controllers	1415

Switch back the MetroCluster configuration	1426
Check the health of the MetroCluster configuration	1428
Upgrade the nodes on cluster_A	1429
Send a custom AutoSupport message after maintenance	1429
Restore Tiebreaker monitoring	1429
Upgrade controllers from AFF A700/FAS9000 to AFF A900/FAS9500 in a MetroCluster FC configuration using switchover and switchback (ONTAP 9.10.1 or later)	1429
Enable console logging	1430
Prepare for the upgrade	1431
Clear slot 7 on the AFF A700 controller	1431
Switch over the MetroCluster configuration	1438
Remove the AFF A700 or FAS9000 controller module and NVS at site_B	1439
Remove the controller module and NVS from both nodes at site_B	1440
Install the AFF A900 or FAS9500 NVS and controller module	1441
Switch back the MetroCluster configuration	1454
Check the health of the MetroCluster configuration	1455
Upgrade the nodes on site_A	1456
Send a custom AutoSupport message after maintenance	1456
Restore Tiebreaker monitoring	1456
Upgrade controllers in a four-node MetroCluster FC configuration using switchover and switchback with "system controller replace" commands (ONTAP 9.10.1 and later)	1457
Supported platform combinations	1457
About this task	1457
Enable console logging	1458
Prepare for the upgrade	1458
Replace the old controllers and boot up the new controllers	1462
Complete the upgrade	1473
Upgrade the nodes on cluster_A	1474
Refreshing a four-node MetroCluster FC configuration	1474
Enable console logging	1474
Perform the refresh procedure	1474
Refresh a four-node or an eight-node MetroCluster IP configuration (ONTAP 9.8 and later)	1476
Enable console logging	1478
Perform the refresh procedure	1478
Expand a two-node MetroCluster FC configuration to a four-node configuration	1487
Expanding a two-node MetroCluster FC configuration to a four-node configuration	1487
Enable console logging	1489
Verifying the state of the MetroCluster configuration	1489
Sending a custom AutoSupport message before adding nodes to the MetroCluster configuration	1491
Zoning for the new controller ports when adding a controller module in a fabric-attached MetroCluster configuration	1491
Add a new controller module to each cluster	1492
Refreshing the MetroCluster configuration with new controllers	1524
Enabling storage failover on both controller modules and enabling cluster HA	1527
Restarting the SVMs	1528

Expand a four-node MetroCluster FC configuration to an eight-node configuration	1528
Expanding a four-node MetroCluster FC configuration to an eight-node configuration	1528
Enable console logging	1531
Determining the new cabling layout	1531
Racking the new equipment	1532
Verifying the health of the MetroCluster configuration	1532
Checking for MetroCluster configuration errors with Config Advisor	1534
Sending a custom AutoSupport message prior to adding nodes to the MetroCluster configuration . . .	1534
Recable and zone a switch fabric for the new nodes	1535
Configure ONTAP on the new controllers	1536
Checking for MetroCluster configuration errors with Config Advisor	1566
Sending a custom AutoSupport message after to adding nodes to the MetroCluster configuration . . .	1566
Verifying switchover, healing, and switchback	1567
Expand a MetroCluster IP configuration	1567
Enable console logging	1569
Example naming in this procedure	1569
Supported platform combinations when adding a second DR group	1569
Sending a custom AutoSupport message prior to maintenance	1572
Considerations for VLANs when adding a new DR group	1573
Verifying the health of the MetroCluster configuration	1574
Removing the configuration from monitoring applications	1578
Preparing the new controller modules	1579
Upgrade RCF files	1579
Join the new nodes to the clusters	1581
Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates . . .	1582
Finalizing the addition of the new nodes	1594
Remove a DR group from a MetroCluster configuration	1600
Enable console logging	1600
Remove the DR group nodes from each cluster	1601
Where to find additional information	1607
MetroCluster and miscellaneous information	1607
Recover from a disaster	1610
Workflow for disaster recovery	1610
Performing a forced switchover after a disaster	1610
Fencing off the disaster site	1610
Performing a forced switchover	1611
Output for the storage aggregate plex show command is indeterminate after a MetroCluster switchover	1612
Accessing volumes in NVFAIL state after a switchover	1612
Choosing the correct recovery procedure	1612
Controller module failure scenarios during MetroCluster installation	1613
Controller module failure scenarios during MetroCluster FC-to-IP transition	1614
Controller module failure scenarios in eight-node MetroCluster configurations	1614
Controller module failure scenarios in two-node MetroCluster configurations	1618
Recover from a multi-controller or storage failure	1619

Recovering from a multi-controller or storage failure	1619
Enable console logging	1620
Replace hardware and boot new controllers	1620
Prepare for switchback in a MetroCluster IP configuration	1632
Prepare for switchback in a MetroCluster FC configuration	1677
Preparing for switchback in a mixed configuration (recovery during transition)	1705
Completing recovery	1708
Recovering from a non-controller failure	1721
Enable console logging	1721
Healing the configuration in a MetroCluster configuration	1722
Verifying that your system is ready for a switchback	1724
Performing a switchback	1726
Verifying a successful switchback	1728
Deleting stale aggregate listings after switchback	1730
Legal notices	1732
Copyright	1732
Trademarks	1732
Patents	1732
Privacy policy	1732
Safety information and regulatory notices	1732

MetroCluster Documentation

MetroCluster release notes

What's new in MetroCluster features

Each release of the ONTAP 9 data management software delivers new and enhanced features that improve the capabilities, manageability, and performance of ONTAP MetroCluster configurations.

For details about known issues, limitations, and upgrade cautions affecting ONTAP MetroCluster configurations, refer to the [ONTAP 9 Release Notes](#). You must sign in with your NetApp account or create an account to access the Release Notes.

Supported features in MetroCluster configuration	Description and where to learn more	Available beginning
SnapMirror cloud support for MetroCluster FlexGroup volumes	SnapMirror cloud supports backup and restore operations for FlexGroup volumes in MetroCluster configurations. Back up data to the cloud using ONTAP SnapMirror	ONTAP 9.18.1GA
New supported upgrade combinations for MetroCluster IP controller upgrades using <code>system controller replace</code> commands	Support for AFF A70 to AFF A90 and FAS70 to FAS90 system upgrades using <code>system controller replace</code> commands in a MetroCluster IP configuration. Upgrade controllers in four-node MetroCluster IP using switchover and switchback with "system controller replace" commands (ONTAP 9.13.1 and later)	ONTAP 9.18.1GA
Flash Cache support for FAS50 systems in MetroCluster IP configurations	Flash Cache is supported on FAS50 systems in MetroCluster IP configurations. Disk assignment on FAS50 systems with Flash Cache	ONTAP 9.18.1
MetroCluster IP support for end-to-end encryption	End-to-end encryption is supported on the following systems to encrypt back-end traffic, such as NVlog and storage replication data, between the sites in a MetroCluster IP configuration. <ul style="list-style-type: none">• AFF A800, AFF C800• AFF A20, AFF A30, AFF C30, AFF A50, AFF C60• AFF A70, AFF A90, AFF A1K, AFF C80• FAS50, FAS70, FAS90 Configure end-to-end encryption in a MetroCluster IP configuration	ONTAP 9.17.1

Supported features in MetroCluster configuration	Description and where to learn more	Available beginning
Limit changes for MetroCluster IP configurations	<p>ONTAP 9.17.1 includes the following limit updates for four-node MetroCluster IP configurations:</p> <ul style="list-style-type: none"> • AFF C800, AFF A800, AFF A900, AFF A90, and AFF A1K systems have the following updated limits: <ul style="list-style-type: none"> ◦ FlexVol volume limits per node: 1250 ◦ SVM limits: 64 SVMs per cluster ◦ LIF count: 256 LIFs per cluster • AFF A400, AFF C400, ASAA400, ASA C400, AFF C80, AFF A70, and AFF A50 systems have the following updated limits: <ul style="list-style-type: none"> ◦ FlexVol volume limits per aggregate (single or multiple): 625 ◦ FlexVol volume limits per node: 1250 ◦ FlexVol volume limits per high-availability (HA) pair: 2500 ◦ FlexVol volume limits per cluster: 5000 ◦ SVM limits: 64 SVMs per cluster ◦ LIF count: 256 LIFs per cluster <p>Refer to the Hardware Universe for more information.</p>	ONTAP 9.17.1
FibreBridge firmware update using credentials	<p>You can update the firmware on FibreBridge bridges using credentials if they are required by the server to download the firmware package.</p> <p>Update firmware on a FibreBridge bridge</p>	ONTAP 9.16.1
SVM data mobility support for migrating MetroCluster configurations	<p>ONTAP supports the following MetroCluster SVM migrations:</p> <ul style="list-style-type: none"> • Migrating an SVM between a non-MetroCluster HA pair and a MetroCluster IP configuration • Migrating an SVM between two MetroCluster IP configurations • Migrating an SVM between a MetroCluster FC configuration and a MetroCluster IP configuration <p>SVM data mobility</p>	ONTAP 9.16.1
MD5 authentication support for BGP peer groups	<p>ONTAP supports MD5 authentication on BGP peer groups to protect BGP sessions. When MD5 is enabled, BGP sessions can only be established and processed among authorized peers, preventing potential disruptions of the session by an unauthorized actor.</p> <p>Configure virtual IP (VIP) LIFs</p>	ONTAP 9.16.1

Supported features in MetroCluster configuration	Description and where to learn more	Available beginning
MetroCluster IP support for end-to-end encryption	<p>End-to-end encryption is supported on AFF A400, AFF C400, FAS8300, and FAS8700 systems to encrypt back-end traffic, such as NVlog and storage replication data, between the sites in a MetroCluster IP configuration.</p> <p>Configure end-to-end encryption in a MetroCluster IP configuration</p>	ONTAP 9.15.1
Volume limit increase for four-node MetroCluster IP configurations on AFF A800 and AFF C800 systems	<p>In four-node MetroCluster IP configurations, the following volume limits for AFF A800 and AFF C800 systems have increased:</p> <ul style="list-style-type: none"> • The maximum number of FlexVol volumes per aggregate increased from 200 to 625. • The maximum number of FlexVol volumes per node increased from 800 to 1250. • The maximum number of FlexVol volumes per HA pair increased from 1600 to 2500. 	ONTAP 9.15.1
MetroCluster IP support for NVMe	<p>The NVMe/TCP front-end host protocol is supported on four-node MetroCluster IP configurations.</p> <p>SAN configurations in a MetroCluster environment</p>	ONTAP 9.15.1
Volume limit increase for four-node MetroCluster IP configurations on AFF A900 systems	<p>In four-node MetroCluster IP configurations, the following volume limits for AFF A900 systems have increased:</p> <ul style="list-style-type: none"> • The maximum number of FlexVol volumes per aggregate increased from 200 to 625. • The maximum number of FlexVol volumes per node increased from 800 to 1250. • The maximum number of FlexVol volumes per HA pair increased from 1600 to 2500. 	ONTAP 9.14.1
S3 object storage support on mirrored and unmirrored aggregates	<p>You can enable an S3 object storage server on an SVM in a mirrored or unmirrored aggregate in MetroCluster IP and FC configurations.</p> <p>S3 support with MetroCluster</p>	ONTAP 9.14.1
Support for provisioning an S3 bucket on mirrored and unmirrored aggregates in a MetroCluster cluster	<p>You can create a bucket on a mirrored or unmirrored aggregate in MetroCluster configurations.</p> <p>Create an ONTAP S3 bucket on a mirrored or unmirrored aggregate in a MetroCluster configuration</p>	ONTAP 9.14.1

Supported features in MetroCluster configuration	Description and where to learn more	Available beginning
Transition from MetroCluster FC to MetroCluster IP using a shared switch for MetroCluster IP and Ethernet attached storage	<p>You can transition nondisruptively from a MetroCluster FC to a MetroCluster IP configuration using a shared storage switch.</p> <p>Transition nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later)</p>	ONTAP 9.13.1
Nondisruptive transitions from an eight-node MetroCluster FC configuration to a MetroCluster IP configuration	<p>You can nondisruptively transition workloads and data from an existing eight-node MetroCluster FC configuration to a new MetroCluster IP configuration.</p> <p>Transition nondisruptively from a MetroCluster FC to a MetroCluster IP configuration</p>	ONTAP 9.13.1
Four-node MetroCluster IP configuration upgrades using switchover and switchback	<p>You can upgrade controllers in a four-node MetroCluster IP configuration using switchover and switchback with <code>system controller replace</code> commands.</p> <p>Upgrade controllers in a four node MetroCluster IP configuration</p>	ONTAP 9.13.1
Mediator-assisted automatic unplanned switchover (MAUSO) is triggered for an environmental shutdown	<p>If one site shuts down gracefully due to an environmental shutdown, MAUSO is triggered.</p> <p>How the ONTAP Mediator supports automatic unplanned switchover</p>	ONTAP 9.13.1
Eight-node MetroCluster IP configurations support	<p>You can upgrade the controllers and storage in an eight-node MetroCluster IP configuration by expanding the configuration to become a temporary twelve-node configuration and then removing the old DR groups.</p> <p>Refresh a four-node MetroCluster IP configuration</p>	ONTAP 9.13.1
MetroCluster IP configuration conversion to a shared storage MetroCluster switch configuration	<p>You can convert a MetroCluster IP configuration to a shared storage MetroCluster switch configuration.</p> <p>Replace an IP switch</p>	ONTAP 9.13.1

Supported features in MetroCluster configuration	Description and where to learn more	Available beginning
MetroCluster automatic forced switchover feature in a MetroCluster IP configuration	<p>You can enable the MetroCluster automatic forced switchover feature in a MetroCluster IP configuration. This feature is an extension of the Mediator-assisted unplanned switchover (MAUSO) feature.</p> <p>Automatic switchover limitations</p>	ONTAP 9.12.1
S3 on an SVM on an unmirrored aggregate in a MetroCluster IP configuration	<p>You can enable an ONTAP Simple Storage Service (S3) object storage server on an SVM on an unmirrored aggregate in a MetroCluster IP configuration.</p> <p>S3 support with MetroCluster</p>	ONTAP 9.12.1
MetroCluster IP support for NVMe	<p>The NVMe/FC protocol is supported on four-node MetroCluster IP configurations.</p> <p>SAN configurations in a MetroCluster environment</p>	ONTAP 9.12.1
IPsec support for front-end host protocol in MetroCluster IP and MetroCluster fabric-attached configurations	<p>IPsec support for front-end host protocol (such as NFS and iSCSI) is available in MetroCluster IP and MetroCluster fabric-attached configurations.</p> <p>Configure IP security (IPsec) over wire encryption</p>	ONTAP 9.12.1
Transition from a MetroCluster FC configuration to an AFF A250 or FAS500f MetroCluster IP configuration	<p>You can transition from a MetroCluster FC configuration to an AFF A250 or FAS500f MetroCluster IP configuration.</p> <p>Move the local cluster connections</p>	ONTAP 9.11.1
Consistency groups	<p>Consistency groups are supported in MetroCluster configurations.</p> <p>Consistency groups in MetroCluster configurations</p>	ONTAP 9.11.1
Simplified controller upgrade of nodes in a MetroCluster FC configuration	<p>The upgrade procedure for the upgrade process using switchover and switchback has been simplified.</p> <p>Upgrade controllers in a MetroCluster FC configuration using switchover and switchback</p>	ONTAP 9.10.1
IP support for shared link at layer 3	<p>MetroCluster IP configurations can be implemented with IP-routed (layer 3) back-end connections.</p> <p>Considerations for layer 3 wide-area networks</p>	ONTAP 9.9.1

Supported features in MetroCluster configuration	Description and where to learn more	Available beginning
Support for eight-node MetroCluster configurations	Permanent eight-node clusters are supported in IP and Fabric-attached MetroCluster configurations. Install and cable MetroCluster components	ONTAP 9.9.1

What's new in MetroCluster IP platform and switch support

Learn what's new in MetroCluster IP platform and switch support.

Platform support

Supported platforms in MetroCluster IP configurations	Available beginning
FAS50	ONTAP 9.16.1GA
AFF A20, AFF A30, AFF A50, AFF C30, AFF C60, AFF C80	ONTAP 9.16.1
FAS70, FAS90	ONTAP 9.15.1P3
AFF A70, AFF A90, AFF A1K	ONTAP 9.15.1
ASAA150, ASAA250, ASA A400, ASA A800, ASA A900, ASA C250, ASA C400, ASA C800	ONTAP 9.14.1
AFF A150	ONTAP 9.13.1 ONTAP 9.12.1P1 ONTAP 9.11.1P8 ONTAP 9.10.1P12
AFF C250, AFF C400, AFF C800	ONTAP 9.12.1P1 ONTAP 9.13.1 GA
AFF A900	ONTAP 9.10.1
AFF A250	ONTAP 9.8
FAS500f	ONTAP 9.8

Supported platforms in MetroCluster IP configurations	Available beginning
ASA AFF A220, ASA AFF A250, ASA AFF A400, ASA AFF A700, ASA AFF A800	ONTAP 9.7
AFF A320	ONTAP 9.6P3
AFF A220, FAS2750	ONTAP 9.6
AFF A300, FAS8200	ONTAP 9.5

Switch support

Broadcom IP switches	Available beginning
BES-53248	ONTAP 9.6

Cisco IP switches	Available beginning
9336C-FX2 (12-port)	ONTAP 9.14.1
9336C-FX2 (36-port)	ONTAP 9.8
3132Q-V	ONTAP 9.6
3232C	ONTAP 9.6

NVIDIA switches	Available beginning
Multiple MetroCluster IP configurations on the same NVIDIA SN2100 switch	ONTAP 9.14.1
SN2100	ONTAP 9.12.1

What's new in MetroCluster FC platform and switch support

Learn what's new in MetroCluster FC platform and switch support.

Platform support

Supported platforms in MetroCluster FC configurations	Available beginning
AFF A900	ONTAP 9.10.1

Supported platforms in MetroCluster FC configurations	Available beginning
ASA AFF A700 and ASA AFF A400	ONTAP 9.7P5
AFF A400 and FAS8300	ONTAP 9.7
AFF A300 and FAS8200	ONTAP 9.5

Switch support

Brocade FC switches	Available beginning
G710	ONTAP 9.17.1
G720	ONTAP 9.8
G620-1, G630-1	ONTAP 9.8
G630	ONTAP 9.6

What's new in ONTAP Mediator support for MetroCluster IP

Learn about the new MetroCluster IP features and enhancements for ONTAP Mediator support.

For details on the features and enhancements for each release of ONTAP Mediator, refer to [What's new in ONTAP Mediator](#).

ONTAP Mediator capability	Available beginning
<p>IPv6 is supported for ONTAP Mediator 1.11 or later in MetroCluster IP configurations.</p> <p>Set up the ONTAP Mediator for a MetroCluster IP configuration</p>	ONTAP 9.18.1
<p>ONTAP Mediator 1.11 adds support for managing up to ten MetroCluster IP configurations using a single ONTAP Mediator instance.</p> <p>Prepare to install the ONTAP Mediator in a MetroCluster IP configuration</p>	ONTAP 9.18.1
<p>Mediator-assisted automatic unplanned switchover (MAUSO) is supported in the case of an environmental shutdown.</p> <p>If one site shuts down gracefully due to an environmental shutdown, MAUSO is triggered.</p> <p>How ONTAP Mediator supports automatic unplanned switchover</p>	ONTAP 9.13.1

ONTAP Mediator capability	Available beginning
Initial support for ONTAP Mediator in MetroCluster IP configurations	ONTAP 9.7

What's new in MetroCluster Tiebreaker support

Enhancements to the MetroCluster Tiebreaker software are provided with each release. Here's what's new in recent releases of MetroCluster Tiebreaker.

Enhancements

ONTAP Tiebreaker version	Enhancements
1.7	<ul style="list-style-type: none"> • Bug fixes • Adds support for switchover simulation using the CLI
1.6P1	<ul style="list-style-type: none"> • Supporting libraries update • Security enhancements
1.6	<ul style="list-style-type: none"> • Improved ease of installation • Supporting libraries update • Security enhancements
1.5	<ul style="list-style-type: none"> • Supporting libraries update • Security enhancements
1.4	<ul style="list-style-type: none"> • Supporting libraries update

OS support matrix

The following table indicates the supported operating systems for each version of Tiebreaker.

OS for Tiebreaker	1.7	1.6P1	1.6	1.5	1.4
Rocky Linux 9.4	Yes	Yes	No	No	No
Rocky Linux 9.0	No	No	Yes	No	No
Rocky Linux 8.10	Yes	Yes	No	No	No

Red Hat Enterprise Linux (RHEL) 9.6	Yes	Yes	No	No	No
RHEL 9.5	Yes	Yes	No	No	No
RHEL 9.4	Yes	Yes	No	No	No
RHEL 9.3	No	No	No	No	No
RHEL 9.2	Yes	Yes	Yes	No	No
RHEL 9.1	No	No	Yes	No	No
RHEL 9.0	No	No	Yes	No	No
RHEL 8.11 - 9.0	No	No	Yes	No	No
RHEL 8.10	Yes	Yes	Yes	No	No
RHEL 8.9	No	No	Yes	No	No
RHEL 8.8	Yes	Yes	Yes	No	No
RHEL 8.1 - 8.7	No	No	Yes	Yes	Yes
RHEL 7 - 7.9	No	No	No	No	Yes
CentOS 7 - 7.9	No	No	No	No	Yes

Install a fabric-attached MetroCluster

Overview

To install your fabric-attached MetroCluster configuration, you must perform a number of procedures in the correct order.

- [Prepare for the installation and understand all requirements.](#)
- [Cable the components](#)
- [Configure the software](#)
- [Test the configuration](#)

Prepare for the MetroCluster installation

Differences among the ONTAP MetroCluster configurations

The various MetroCluster configurations have key differences in the required components.

In all configurations, each of the two MetroCluster sites are configured as an ONTAP cluster. In a two-node MetroCluster configuration, each node is configured as a single-node cluster.

Feature	IP configurations	Fabric attached configurations		Stretch configurations	
		Four- or eight-node	Two-node	Two-node bridge-attached	Two-node direct-attached
Number of controllers	Four or eight ¹	Four or eight	Two	Two	Two
Uses an FC switch storage fabric	No	Yes	Yes	No	No
Uses an IP switch storage fabric	Yes	No	No	No	No
Uses FC-to-SAS bridges	No	Yes	Yes	Yes	No
Uses direct-attached SAS storage	Yes (local attached only)	No	No	No	Yes

Supports ADP	Yes (beginning with ONTAP 9.4)	No	No	No	No
Supports local HA	Yes	Yes	No	No	No
Supports ONTAP automatic unplanned switchover (AUSO)	No	Yes	Yes	Yes	Yes
Supports unmirrored aggregates	Yes (beginning with ONTAP 9.8)	Yes	Yes	Yes	Yes
Supports ONTAP Mediator	Yes (beginning with ONTAP 9.7)	No	No	No	No
Supports MetroCluster Tiebreaker	Yes (not in combination with ONTAP Mediator)	Yes	Yes	Yes	Yes
Supports All SAN Arrays	Yes	Yes	Yes	Yes	Yes

Notes

- Review the following considerations for eight-node MetroCluster IP configurations:
 - Eight-node configurations are supported beginning with ONTAP 9.9.1.
 - Only NetApp-validated MetroCluster switches (ordered from NetApp) are supported.
 - Configurations using IP-routed (layer 3) backend connections are not supported.

Support for All SAN Array systems in MetroCluster configurations

Some of the All SAN Arrays (ASAs) are supported in MetroCluster configurations. In the MetroCluster documentation, the information for AFF models applies to the corresponding ASA system. For example, all cabling and other information for the AFF A400 system also applies to the ASA AFF A400 system.

Supported platform configurations are listed in the [NetApp Hardware Universe](#).

Cluster peering

Each MetroCluster site is configured as a peer to its partner site. You must be familiar with the prerequisites and guidelines for configuring the peering relationships. This is important when deciding on whether to use shared or dedicated ports for those relationships.

Related information

[Cluster and SVM peering express configuration](#)

Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that connectivity between port, IP address, subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node.

For example, in a four-node cluster, the subnet used for intercluster communication must have four available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.



ONTAP 9 enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

- All ports used to communicate with a given remote cluster must be in the same IPspace.

You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.

- The broadcast domain used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).

- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
- TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
- Bidirectional HTTPS between the intercluster LIFs

The default intercluster firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Considerations when using dedicated ports

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether using a dedicated port is the best intercluster network solution:

- If the amount of available WAN bandwidth is similar to that of the LAN ports, and the replication interval is such that replication occurs while regular client activity exists, then you should dedicate Ethernet ports for intercluster replication to avoid contention between replication and the data protocols.
- If the network utilization generated by the data protocols (CIFS, NFS, and iSCSI) is such that the network utilization is above 50 percent, then dedicate ports for replication to allow for nondegraded performance if a node failover occurs.
- When physical 10 GbE or faster ports are used for data and replication, you can create VLAN ports for replication and dedicate the logical ports for intercluster replication.

The bandwidth of the port is shared between all VLANs and the base port.

- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.

Considerations when sharing data ports

When determining whether sharing a data port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether sharing data ports is the best intercluster connectivity solution:

- For a high-speed network, such as a 40-Gigabit Ethernet (40-GbE) network, a sufficient amount of local LAN bandwidth might be available to perform replication on the same 40-GbE ports that are used for data access.

In many cases, the available WAN bandwidth is far less than the 10 GbE LAN bandwidth.

- All nodes in the cluster might have to replicate data and share the available WAN bandwidth, making data port sharing more acceptable.
- Sharing ports for data and replication eliminates the extra port counts required to dedicate ports for replication.
- The maximum transmission unit (MTU) size of the replication network will be the same size as that used on the data network.

- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.
- When data ports for intercluster replication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node to control the specific data port that is used for replication.

Considerations for ISLs

You must determine how many ISLs you need for each FC switch fabric in the MetroCluster configuration. Beginning with ONTAP 9.2, in some cases, instead of dedicating FC switches and ISLs to each individual MetroCluster configuration, you can share the same four switches.

ISL sharing considerations (ONTAP 9.2)

Beginning with ONTAP 9.2, you can use ISL sharing in the following cases:

- One two-node and one four-node MetroCluster configurations
- Two separate four-node MetroCluster configurations
- Two separate two-node MetroCluster configurations
- Two DR groups within one eight-node MetroCluster configuration

The number of ISLs required between the shared switches depends on the bandwidth of the platform models connected to the shared switches.

Consider the following aspects of your configuration when determining how many ISLs you need.

- Non-MetroCluster devices should not be connected to any of the FC switches that provide the back-end MetroCluster connectivity.
- ISL sharing is supported on all switches except the Cisco 9250i and Cisco 9148 switches.
- All nodes must be running ONTAP 9.2 or later.
- The FC switch cabling for ISL sharing is the same as for the eight-node MetroCluster cabling.
- The RCF files for ISL sharing are same as for the eight-node MetroCluster cabling.
- You should verify that all hardware and software versions are supported.

[NetApp Hardware Universe](#)

- The speed and number of ISLs must be sized to support the client load on both MetroCluster systems.
- The back-end ISLs and the back-end components must be dedicated to the MetroCluster configuration only.
- The ISL must use one of the supported speeds: 4 Gbps, 8 Gbps, 16 Gbps, or 32 Gbps.
- The ISLs on one fabric should all be the same speed and length.
- The ISLs on one fabric should all have the same topology. For example, they should all be direct links, or if your system uses WDM, then they should all use WDM.

Platform-specific ISL considerations

The number of recommended ISLs is platform-model specific. The following table shows the ISL requirements for each fabric by platform model. It assumes that each ISL has a 16-Gbps capacity.

Platform model	Recommended number of ISLs per four-node DR group (per switch fabric)
AFF A900 and FAS9500	Eight
AFF A700	Six
FAS9000	Six
8080	Four
All others	Two

If the switch fabric is supporting eight nodes (either part of a single, eight-node MetroCluster configuration, or two four-node configurations that are sharing ISLs), the recommended total number of ISLs for the fabric is the sum of that required for each four-node DR group. For example:

- If DR group 1 includes four AFF A700 systems, it requires six ISLs.
- If DR group 2 includes four FAS8200 systems, it requires two ISLs.
- The total number of recommended ISLs for the switch fabric is eight.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

The Hardware Universe tool provides some notes about the requirements that Time Division Multiplexing (TDM) or Wavelength Division Multiplexing (WDM) equipment must meet to work with a fabric-attached MetroCluster configuration. These notes also include information about various configurations, which can help you to determine when to use in-order delivery (IOD) of frames or out-of-order delivery (OOD) of frames.

An example of such requirements is that the TDM/WDM equipment must support the link aggregation (trunking) feature with routing policies. The order of delivery (IOD or OOD) of frames is maintained within a switch, and is determined by the routing policy that is in effect.

[NetApp Hardware Universe](#)

The following table provides the routing policies for configurations containing Brocade switches and Cisco switches:

Switches	Configuring MetroCluster configurations for IOD	Configuring MetroCluster configurations for OOD
----------	---	---

Brocade	<ul style="list-style-type: none"> • AptPolicy must be set to 1 • DLS must be set to off • IOD must be set to on 	<ul style="list-style-type: none"> • AptPolicy must be set to 3 • DLS must be set to on • IOD must be set to off
Cisco	<p>Policies for the FCVI-designated VSAN:</p> <ul style="list-style-type: none"> • Load balancing policy: srcid and dstid • IOD must be set to on <p>Policies for the storage-designated VSAN:</p> <ul style="list-style-type: none"> • Load balancing policy: srcid, dstid, and oxid • VSAN must not have the in-order-guarantee option set 	Not applicable

When to use IOD

It is best to use IOD if it is supported by the links. The following configurations support IOD:

- A single ISL
- The ISL and the link (and the link equipment, such as TDM/WDM, if used) supports configuration for IOD.
- A single trunk, and the ISLs and the links (and the link equipment, such as TDM/WDM, if used) support configuration for IOD.

When to use OOD

- You can use OOD for all configurations that do not support IOD.
- You can use OOD for configurations that do not support the trunking feature.

Using encryption devices

When using dedicated encryption devices on the ISL or encryption on WDM devices in the MetroCluster configuration, you must meet the following requirements:

- The external encryption devices or WDM equipment has been self certified by the vendor with the FC switch in question.

The self certification should cover the operating mode (such as trunking and encryption).

- The added latency due to encryption should be no more than 10 microseconds.

Requirements for using a Brocade DCX 8510-8 switch

As you prepare for the MetroCluster installation, you should understand the MetroCluster hardware architecture and required components.

- The DCX 8510-8 switches used in MetroCluster configurations must be purchased from NetApp.
- For scalability, you should leave one port-chunk between MetroCluster configurations if cabling only two MetroClusters in 4x48-port modules. This enables you to expand port usage in MetroCluster configurations without recabling.
- Each Brocade DCX 8510-8 switch in the MetroCluster configuration must be correctly configured for the ISL ports and storage connections. For port usage, see the following section: [Port assignments for FC switches](#).
- ISLs cannot be shared and each MetroCluster requires two ISLs for each fabric.
- The DCX 8510-8 switch used for backend MetroCluster connectivity should not be used for any other connectivity.

Non-MetroCluster devices should not be connected to these switches and non-MetroCluster traffic should not flow through DCX 8510-8 switches.

- One line card can either be connected to ONTAP MetroClusters **or** ONTAP 7-Mode MetroClusters.



RCF files are not available for this switch.

The following are the requirements for using two Brocade DCX 8510-8 switches:

- You must have one DCX 8510-8 switch at each site.
- You must use a minimum of two 48-port blades that contain 16Gb SFPs in each switch.

The following are the requirements for using four DCX 8510-8 switches at each site in a MetroCluster configuration:

- You must have two DCX 8510-8 switches at each site.
- You must use at least one 48-port blade for each DCX 8510-8 switch.
- Each blade is configured as a virtual switch using virtual fabrics.

The following NetApp products are not supported by Brocade DCX 8510-8 switches:

- Config Advisor
- Fabric Health Monitor
- MyAutoSupport (system risks might show false positives)
- Active IQ Unified Manager (formerly OnCommand Unified Manager)



Ensure that all the components needed for this configuration are in the [NetApp Interoperability Matrix Tool](#). Read the notes section in the Interoperability Matrix Tool for information on supported configurations.

Considerations when using unmirrored aggregates

Considerations when using unmirrored aggregates

If your configuration includes unmirrored aggregates, you must be aware of potential access issues that follow switchover operations.

Considerations for unmirrored aggregates when doing maintenance requiring power shutdown

If you are performing a negotiated switchover for maintenance reasons requiring site-wide power shutdown, you should first manually take offline any unmirrored aggregates owned by the disaster site.

If you do not take any unmirrored aggregates offline, nodes at the surviving site might go down due to multi-disk panics. This could occur if switched over unmirrored aggregates go offline, or are missing, because of the loss of connectivity to storage at the disaster site. This is the result of a power shutdown or a loss of ISLs.

Considerations for unmirrored aggregates and hierarchical namespaces

If you are using hierarchical namespaces, you should configure the junction path so that all of the volumes in that path are either on mirrored aggregates only or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates in the junction path might prevent access to the unmirrored aggregates after the switchover operation.

Considerations for unmirrored aggregates and CRS metadata volume and data SVM root volumes

The configuration replication service (CRS) metadata volume and data SVM root volumes must be on a mirrored aggregate. You cannot move these volumes to an unmirrored aggregate. If they are on an unmirrored aggregate, negotiated switchover and switchback operations are vetoed. The MetroCluster check command provides a warning if this is the case.

Considerations for unmirrored aggregates and SVMs

SVMs should be configured on mirrored aggregates only, or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates can result in a switchover operation that exceeds 120 seconds and result in a data outage if the unmirrored aggregates do not come online.

Considerations for unmirrored aggregates and SAN

In ONTAP versions prior to 9.9.1, a LUN should not be located on an unmirrored aggregate. Configuring a LUN on an unmirrored aggregate can result in a switchover operation that exceeds 120 seconds and a data outage.

Firewall usage at MetroCluster sites

Considerations for firewall usage at MetroCluster sites

If you are using a firewall at a MetroCluster site, you must ensure access for required ports.

The following table shows TCP/UDP port usage in an external firewall positioned between two MetroCluster sites.

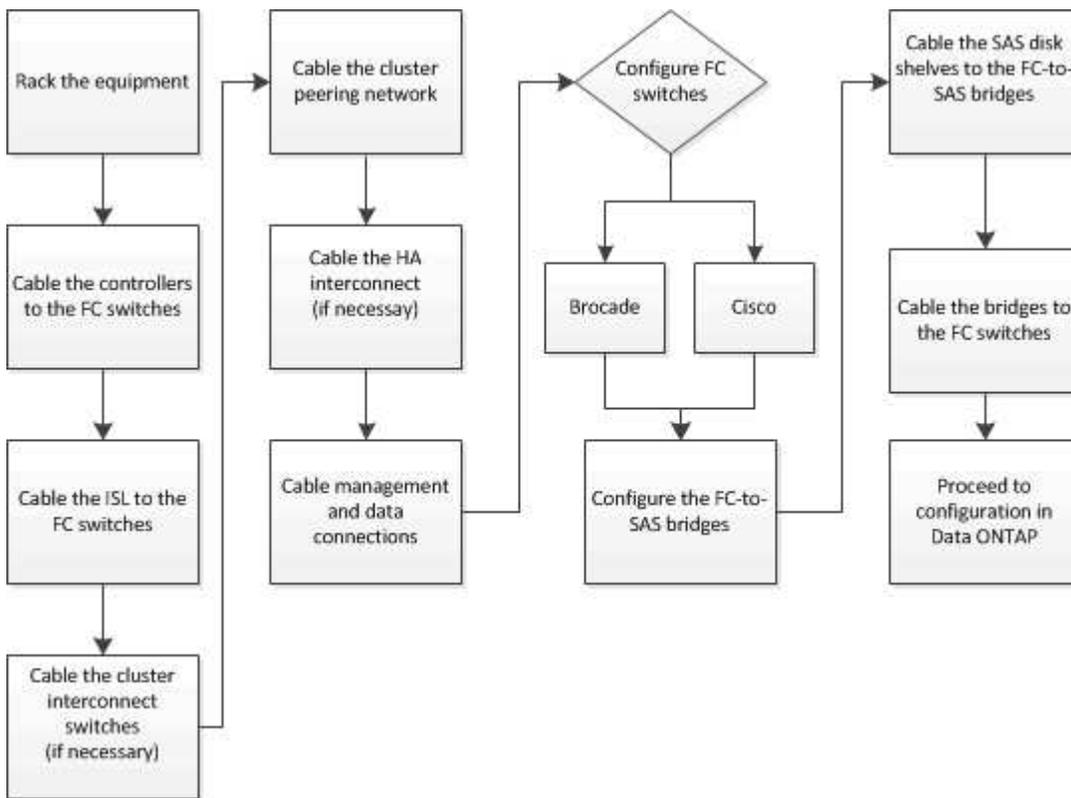
Traffic type	Port/services
Cluster peering	11104 / TCP
	11105 / TCP
ONTAP System Manager	443 / TCP

MetroCluster IP intercluster LIFs	65200 / TCP 10006 / TCP and UDP
Hardware assist	4444 / TCP

Cable a fabric-attached MetroCluster configuration

Cabling a fabric-attached MetroCluster configuration

The MetroCluster components must be physically installed, cabled, and configured at both geographic sites.



Parts of a fabric MetroCluster configuration

Parts of a fabric MetroCluster configuration

As you plan your MetroCluster configuration, you should understand the hardware components and how they interconnect.

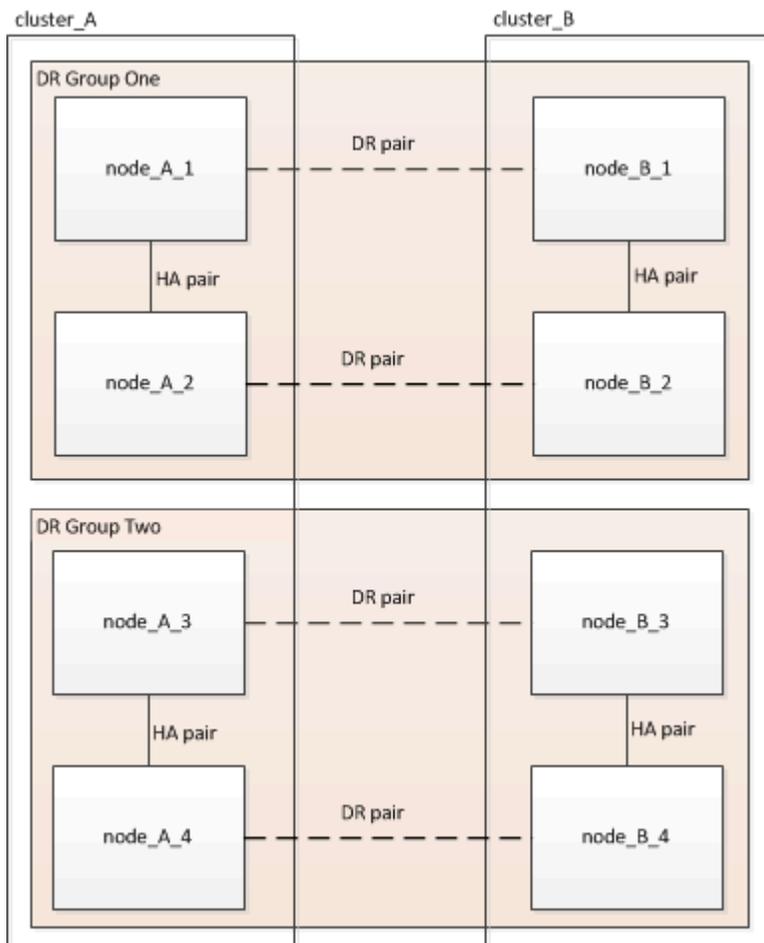
Disaster Recovery (DR) groups

A fabric MetroCluster configuration consists of one or two DR groups, depending on the number of nodes in the MetroCluster configuration. Each DR group consists of four nodes.

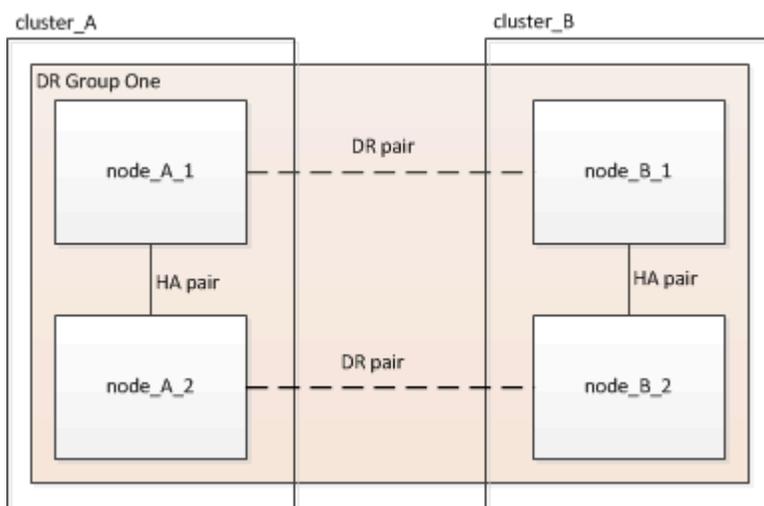
- An eight-node MetroCluster configuration consists of two DR groups.

- A four-node MetroCluster configuration consists of one DR group.

The following illustration shows the organization of nodes in an eight-node MetroCluster configuration:



The following illustration shows the organization of nodes in a four-node MetroCluster configuration:



Key hardware elements

A MetroCluster configuration includes the following key hardware elements:

- Storage controllers

The storage controllers are not connected directly to the storage but connect to two redundant FC switch fabrics.

- FC-to-SAS bridges

The FC-to-SAS bridges connect the SAS storage stacks to the FC switches, providing bridging between the two protocols.

- FC switches

The FC switches provide the long-haul backbone ISL between the two sites. The FC switches provide the two storage fabrics that allow data mirroring to the remote storage pools.

- Cluster peering network

The cluster peering network provides connectivity for mirroring of the cluster configuration, which includes storage virtual machine (SVM) configuration. The configuration of all of the SVMs on one cluster is mirrored to the partner cluster.

Eight-node fabric MetroCluster configuration

An eight-node configuration consists of two clusters, one at each geographically separated site. cluster_A is located at the first MetroCluster site. cluster_B is located at the second MetroCluster site. Each site has one SAS storage stack. Additional storage stacks are supported, but only one is shown at each site. The HA pairs are configured as switchless clusters, without cluster interconnect switches. A switched configuration is supported, but is not shown.

An eight-node configuration includes the following connections:

- FC connections from each controller's HBAs and FC-VI adapters to each of the FC switches
- An FC connection from each FC-to-SAS bridge to an FC switch
- SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
- An HA interconnect between each controller in the local HA pair

If the controllers support a single-chassis HA pair, the HA interconnect is internal, occurring through the backplane, meaning that an external interconnect is not required.

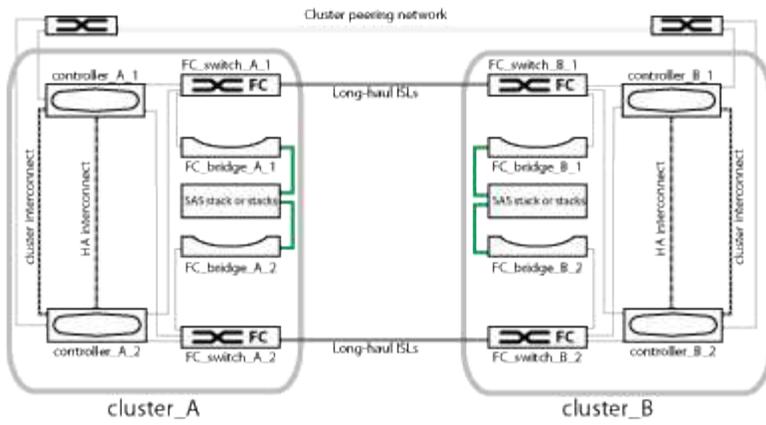
- Ethernet connections from the controllers to the customer-provided network that is used for cluster peering

SVM configuration is replicated over the cluster peering network.

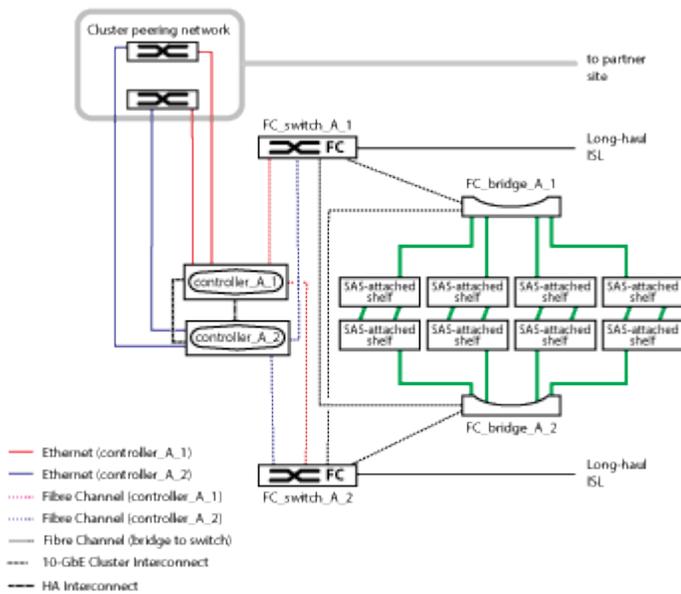
- A cluster interconnect between each controller in the local cluster

Four-node fabric MetroCluster configuration

The following illustration shows a simplified view of a four-node fabric MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

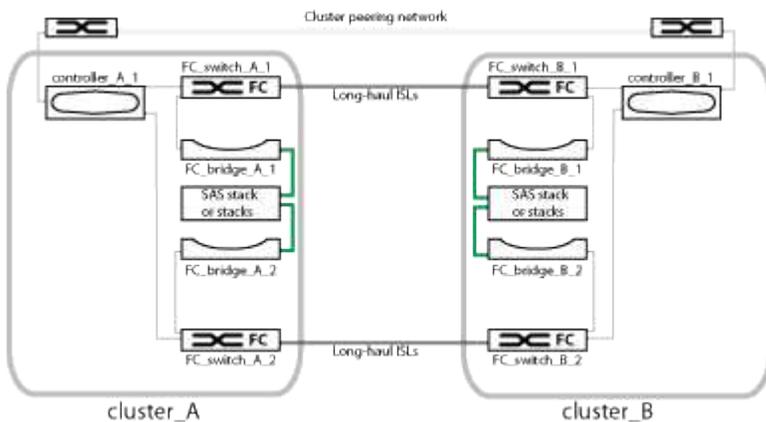


The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):



Two-node fabric MetroCluster configuration

The following illustration shows a simplified view of a two-node fabric MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.

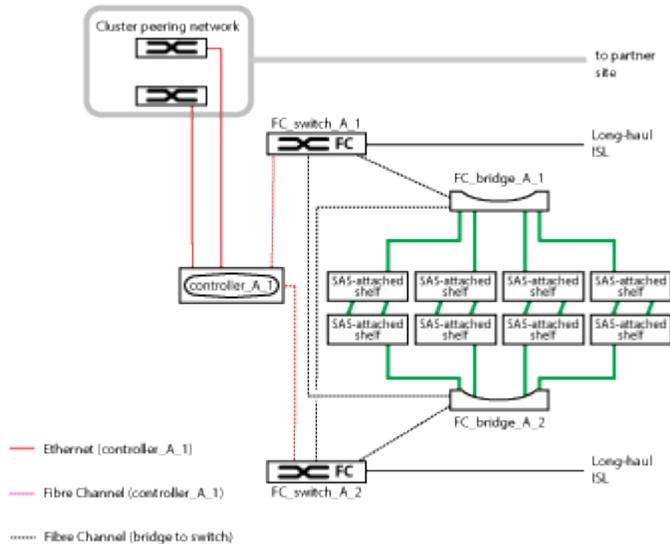


A two-node configuration consists of two clusters, one at each geographically separated site. cluster_A is located at the first MetroCluster site. cluster_B is located at the second MetroCluster site. Each site has one SAS storage stack. Additional storage stacks are supported, but only one is shown at each site.



In a two-node configuration, the nodes are not configured as an HA pair.

The following illustration shows a more detailed view of the connectivity in a single MetroCluster cluster (both clusters have the same configuration):



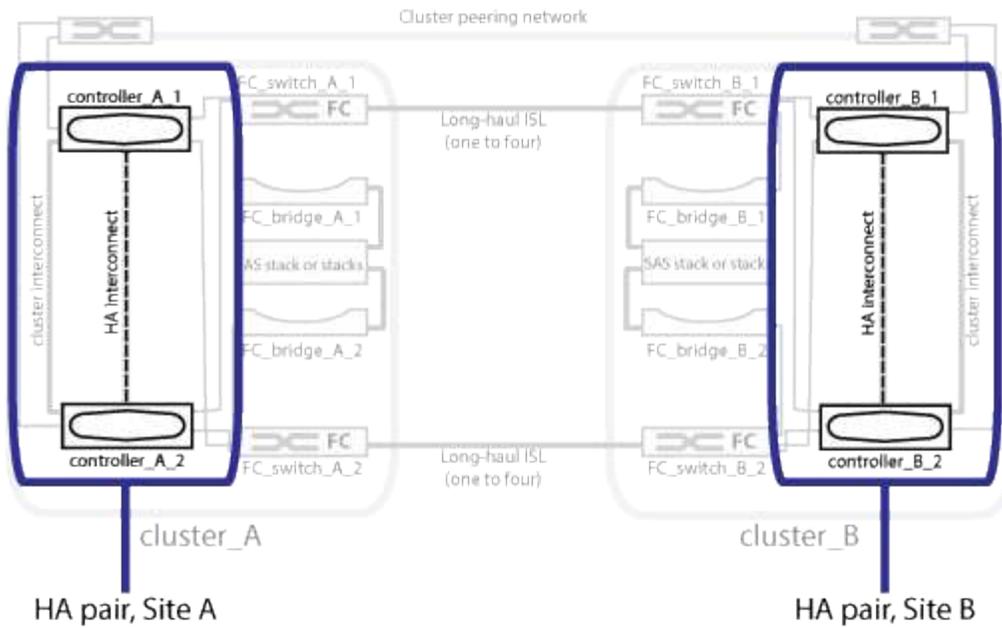
A two-node configuration includes the following connections:

- FC connections between the FC-VI adapter on each controller module
 - FC connections from each controller module's HBAs to the FC-to-SAS bridge for each SAS shelf stack
 - SAS connections between each SAS shelf and from the top and bottom of each stack to an FC-to-SAS bridge
 - Ethernet connections from the controllers to the customer-provided network that is used for cluster peering
- SVM configuration is replicated over the cluster peering network.

Illustration of the local HA pairs in a MetroCluster configuration

In eight-node or four-node MetroCluster configurations, each site consists of storage controllers configured as one or two HA pairs. This allows local redundancy so that if one storage controller fails, its local HA partner can take over. Such failures can be handled without a MetroCluster switchover operation.

Local HA failover and giveback operations are performed with the storage failover commands, in the same manner as a non-MetroCluster configuration.



Related information

[Illustration of redundant FC-to-SAS bridges](#)

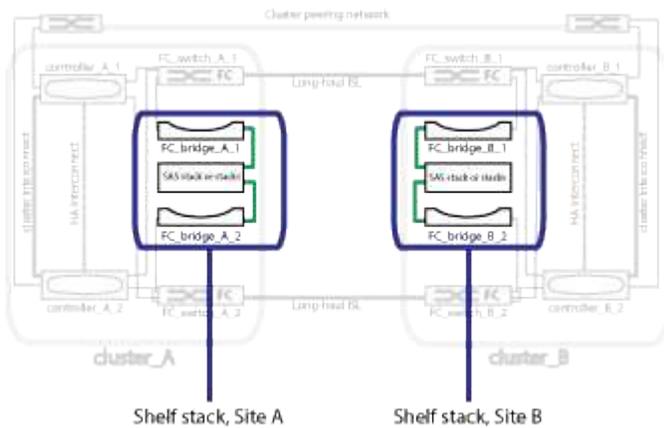
[Redundant FC switch fabrics](#)

[Illustration of the cluster peering network](#)

[ONTAP concepts](#)

Illustration of redundant FC-to-SAS bridges

FC-to-SAS bridges provide protocol bridging between SAS attached disks and the FC switch fabric.



Related information

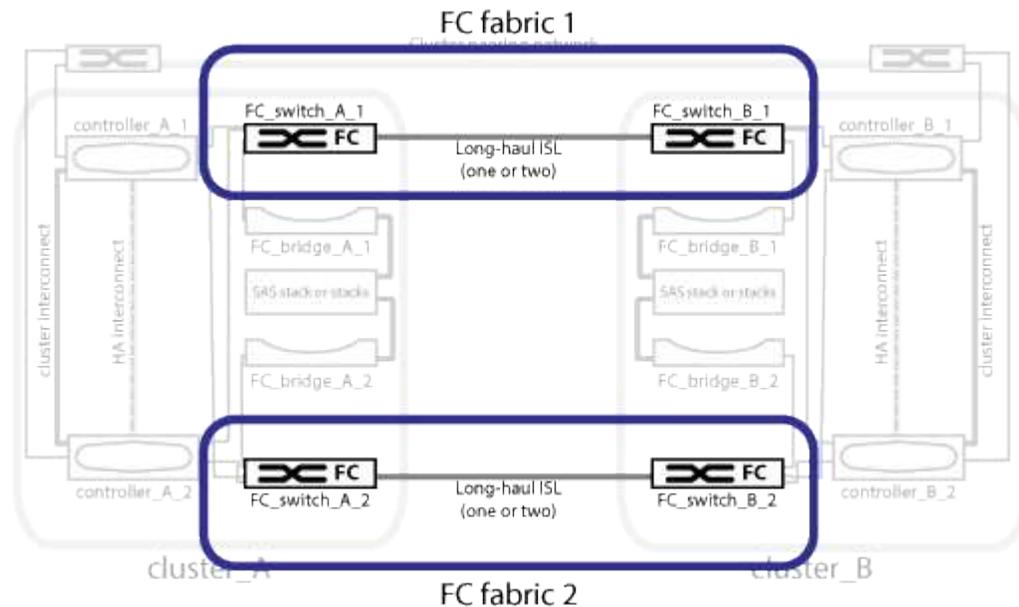
[Illustration of the local HA pairs in a MetroCluster configuration](#)

[Redundant FC switch fabrics](#)

[Illustration of the cluster peering network](#)

Redundant FC switch fabrics

Each switch fabric includes inter-switch links (ISLs) that connect the sites. Data is replicated from site-to-site over the ISL. Each switch fabric must be on different physical paths for redundancy.



Related information

[Illustration of the local HA pairs in a MetroCluster configuration](#)

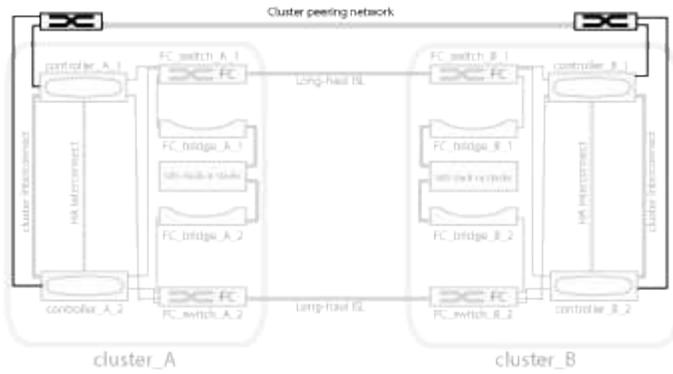
[Illustration of redundant FC-to-SAS bridges](#)

[Illustration of the cluster peering network](#)

Illustration of the cluster peering network

The two clusters in the MetroCluster configuration are peered through a customer-provided cluster peering network. Cluster peering supports the synchronous mirroring of storage virtual machines (SVMs, formerly known as Vservers) between the sites.

Intercluster LIFs must be configured on each node in the MetroCluster configuration, and the clusters must be configured for peering. The ports with the intercluster LIFs are connected to the customer-provided cluster peering network. Replication of the SVM configuration is carried out over this network through the Configuration Replication Service.



Related information

[Illustration of the local HA pairs in a MetroCluster configuration](#)

[Illustration of redundant FC-to-SAS bridges](#)

[Redundant FC switch fabrics](#)

[Cluster and SVM peering express configuration](#)

[Considerations for configuring cluster peering](#)

[Cabling the cluster peering connections](#)

[Peering the clusters](#)

Required MetroCluster FC components and naming conventions

When planning your MetroCluster FC configuration, you must understand the required and supported hardware and software components. For convenience and clarity, you should also understand the naming conventions used for components in examples throughout the documentation. For example, one site is referred to as Site A and the other site is referred to as Site B.

Supported software and hardware

The hardware and software must be supported for the MetroCluster FC configuration.

[NetApp Hardware Universe](#)

When using AFF systems, all controller modules in the MetroCluster configuration must be configured as AFF systems.



Long-wave SFPs are not supported in the MetroCluster storage switches. For a table of supported SFPs, see the MetroCluster Technical Report.

Hardware redundancy in the MetroCluster FC configuration

Because of the hardware redundancy in the MetroCluster FC configuration, there are two of each component at each site. The sites are arbitrarily assigned the letters A and B and the individual components are arbitrarily assigned the numbers 1 and 2.

Requirement for two ONTAP clusters

The fabric-attached MetroCluster FC configuration requires two ONTAP clusters, one at each MetroCluster site.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A
- Site B: cluster_B

Requirement for four FC switches

The fabric-attached MetroCluster FC configuration requires four FC switches (supported Brocade or Cisco models).

The four switches form two switch storage fabrics that provide the ISL between each of the clusters in the MetroCluster FC configuration.

Naming must be unique within the MetroCluster configuration.

Requirement for two, four, or eight controller modules

The fabric-attached MetroCluster FC configuration requires two, four, or eight controller modules.

In a four or eight-node MetroCluster configuration, the controller modules at each site form one or two HA pairs. Each controller module has a DR partner at the other site.

The controller modules must meet the following requirements:

- Naming must be unique within the MetroCluster configuration.
- All controller modules in the MetroCluster configuration must be running the same version of ONTAP.
- All controller modules in a DR group must be of the same model.

However, in configurations with two DR groups, each DR group can consist of different controller module models.

- All controller modules in a DR group must use the same FC-VI configuration.

Some controller modules support two options for FC-VI connectivity:

- Onboard FC-VI ports
- An FC-VI card in slot 1

A mix of one controller module using onboard FC-VI ports and another using an add-on FC-VI card is not supported. For example, if one node uses onboard FC-VI configuration, then all other nodes in the DR group must use onboard FC-VI configuration as well.

Example names:

- Site A: controller_A_1
- Site B: controller_B_1

Requirement for four cluster interconnect switches

The fabric-attached MetroCluster FC configuration requires four cluster interconnect switches (if you are not using two-node switchless clusters)

These switches provide cluster communication among the controller modules in each cluster. The switches are not required if the controller modules at each site are configured as a two-node switchless cluster.

Requirement for FC-to-SAS bridges

The fabric-attached MetroCluster FC configuration requires one pair of FC-to-SAS bridges for each stack group of SAS shelves.



FibreBridge 6500N bridges are not supported in configurations running ONTAP 9.8 and later.

- FibreBridge 7600N or 7500N bridges support up to four SAS stacks.
- Each stack can use different models of IOM.
- Naming must be unique within the MetroCluster configuration.

The suggested names used as examples in this documentation identify the controller module and stack that the bridge connects to, as shown below.

Pool and drive requirements (minimum supported)

Eight SAS disk shelves are recommended (four shelves at each site) to allow disk ownership on a per-shelf basis.

The MetroCluster configuration requires the minimum configuration at each site:

- Each node has at least one local pool and one remote pool at the site.

For example, in a four-node MetroCluster configuration with two nodes at each site, four pools are required at each site.

- At least seven drives in each pool.

In a four-node MetroCluster configuration with a single mirrored data aggregate per node, the minimum configuration requires 24 disks at the site.

In a minimum supported configuration, each pool has the following drive layout:

- Three root drives
- Three data drives
- One spare drive

In a minimum supported configuration, at least one shelf is needed per site.

MetroCluster configurations support RAID-DP and RAID4.

Drive location considerations for partially populated shelves

For correct auto-assignment of drives when using shelves that are half populated (12 drives in a 24-drive shelf), drives should be located in slots 0-5 and 18-23.

In a configuration with a partially populated shelf, the drives must be evenly distributed in the four quadrants of the shelf.

Bridge naming conventions

The bridges use the following example naming:

```
bridge_site_stack_grouplocation in pair
```

This portion of the name...	Identifies the...	Possible values...
site	Site on which the bridge pair physically resides.	A or B
stack group	Number of the stack group to which the bridge pair connects. FibreBridge 7600N or 7500N bridges support up to four stacks in the stack group. The stack group can contain no more than 10 storage shelves.	1, 2, etc.
location in pair	Bridge within the bridge pair. A pair of bridges connect to a specific stack group.	a or b

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

Configuration worksheets for FC switches and FC-to-SAS bridges

Before beginning to configure the MetroCluster sites, you can use the following worksheets to record your site information:

[Site A worksheet](#)

[Site B worksheet](#)

Install and cable MetroCluster components

Racking the hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

About this task

This task must be performed on both MetroCluster sites.

Steps

1. Plan out the positioning of the MetroCluster components.

The rack space depends on the platform model of the controller modules, the switch types, and the number of disk shelf stacks in your configuration.

2. Properly ground yourself.
3. Install the controller modules in the rack or cabinet.

[ONTAP Hardware Systems Documentation](#)

4. Install the FC switches in the rack or cabinet.
5. Install the disk shelves, power them on, and then set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites).
6. Install each FC-to-SAS bridge:

- a. Secure the “L” brackets on the front of the bridge to the front of the rack (flush-mount) with the four screws.

The openings in the bridge “L” brackets are compliant with rack standard ETA-310-X for 19-inch (482.6 mm) racks.

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model contains more information and an illustration of the installation.



For adequate port space access and FRU serviceability, you must leave 1U space below the bridge pair and cover this space with a tool-less blanking panel.

- b. Connect each bridge to a power source that provides a proper ground.
- c. Power on each bridge.



For maximum resiliency, bridges that are attached to the same stack of disk shelves must be connected to different power sources.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed its power-on self test sequence.

Cabling the new controller module's FC-VI and HBA ports to the FC switches

The FC-VI ports and HBAs (host bus adapters) must be cabled to the site FC switches on each controller module in the MetroCluster configuration.

Steps

1. Cable the FC-VI ports and HBA ports, using the table for your configuration and switch model.

- [Port assignments for FC switches](#)

Cabling the ISLs between MetroCluster sites

You must connect the FC switches at each site through the fiber-optic Inter-Switch Links (ISLs) to form the switch fabrics that connect the MetroCluster components.

About this task

This must be done for both switch fabrics.

Steps

1. Connect the FC switches at each site to all ISLs, using the cabling in the table that corresponds to your configuration and switch model.
 - [Port assignments for FC switches](#)

Related information

[Considerations for ISLs](#)

Port assignments for FC switches

Port assignments for MetroCluster FC switches

You need to verify that you are using the specified port assignments when you cable the FC switches.

You can reconfigure ports that are not used for attaching initiator ports, FC-VI ports, or ISLs to act as storage ports. However, if you are using the supported RCFs, you must change the zoning accordingly.

If you use the supported RCFs, ISL ports might not connect to the same ports shown and might need to be reconfigured manually.

If you configured your switches using the port assignments for ONTAP 9, you can continue to use the older assignments. However, new configurations running ONTAP 9.1 or later should use the port assignments shown here.

Overall cabling guidelines

You should be aware of the following guidelines when using the cabling tables:

- The Brocade and Cisco switches use different port numbering:
 - On Brocade switches, the first port is numbered 0.
 - On Cisco switches, the first port is numbered 1.
- The cabling is the same for each FC switch in the switch fabric.
- You can order AFF A300 and FAS8200 storage systems with one of two options for FC-VI connectivity:
 - Onboard ports 0e and 0f configured in FC-VI mode.
 - Ports 1a and 1b on an FC-VI card in slot 1.
- AFF A700 and FAS9000 storage systems require four FC-VI ports. The following tables show cabling for the FC switches with four FC-VI ports on each controller except for the Cisco 9250i switch.

For other storage systems, use the cabling shown in the tables but ignore the cabling for FC-VI ports c and d.

You can leave those ports empty.

- AFF A400 and FAS8300 storage systems use ports 2a and 2b for FC-VI connectivity.
- If you have two MetroCluster configurations sharing ISLs, use the same port assignments as that for an eight-node MetroCluster cabling.
- The number of ISLs you cable can vary depending on site's requirements.
- See the section on ISL considerations.

Considerations for ISLs

AFF A900 and FAS9500 cabling guidelines

- AFF A900 or FAS9500 storage systems require eight FC-VI ports. If you are using an AFF A900 or FAS9500, you need to use the eight port configuration. If the configuration includes the other storage system models, use the cabling shown in the tables but ignore the cabling for unneeded FC-VI ports.

Port assignments for systems using two initiator ports

You can configure FAS8200 and AFF A300 systems using a single initiator port for each fabric and two initiator ports for each controller.

You can follow the cabling for the FibreBridge 7600N bridge using only one FC port (FC1 or FC2). Instead of using four initiators, connect only two initiators and leave the other two that are connected to the switch port empty.

If zoning is performed manually, then follow the zoning used for a FibreBridge 7600N bridge using one FC port (FC1 or FC2). In this scenario, one initiator port rather than two is added to each zone member per fabric.

You can change the zoning or perform an upgrade from a FibreBridge 6500N to a FibreBridge 7500N using the procedure in [Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#).

The following table shows port assignments for Brocade FC switches when using a single initiator port for each fabric and two initiator ports for each controller.

Configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only			
MetroCluster 1 or DR Group 1			
Component	Port	6505, 6510, 6520, 7840, G620, G630, G610, G710, G720, G730 and DCX 8510-8	
		Connects to FC switch...	Connects to switch port...

Configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only			
controller_x_1	FC-VI port a	1	0
	FC-VI port b	2	0
	FC-VI port c	1	1
	FC-VI port d	2	1
	HBA port a	1	2
	HBA port b	2	2
	HBA port c	-	-
	HBA port d	-	-
Stack 1	bridge_x_1a	1	8
	bridge_x_1b	2	8
Stack y	bridge_x_ya	1	11
	bridge_x_yb	2	11

Brocade port usage for controllers in a MetroCluster FC configuration

Learn about the port assignments required to cable Brocade FC switches to your controllers.

The following tables show the maximum supported configuration, with four controller modules per DR group. For smaller configurations, ignore the rows for the additional controller modules. Note that eight ISLs are supported only on the Brocade 6510, Brocade DCX 8510-8, G620, G630, G620-1, G630-1, G720, and G730 switches.

Review the following information before using these tables:

- Port usage for the Brocade 6505, G610, and G710 switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, port assignments must be made on a site-by-site basis depending on the controller module model and the number of ISLs and bridge pairs in use.
- The Brocade DCX 8510-8 switch can use the same port layout as the 6510 switch **or** the 7840 switch.
- Brocade 6520, 7810, and 7840 switches aren't supported on systems that use eight FC-VI ports (AFF A900 and FAS9500 systems).
- Brocade 7810 switches only support one DR group.

MetroCluster 1 or DR group 1

The following table shows the supported controller configurations in MetroCluster 1 or DR group 1 on Brocade switches.

Component	Port	Connect s to FC switch...	6505, G610, G710 port	6510, DCX, 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller_x_1	FC-VI port a	1	0	0	0	0	0	0	0
	FC-VI port b	2	0	0	0	0	0	0	0
	FC-VI port c	1	1	1	1	1	1	1	1
	FC-VI port d	2	1	1	1	1	1	1	1
	FC-VI-2 port a	1	16	20	n/a	n/a	n/a	16	2
	FC-VI-2 port b	2	16	20	n/a	n/a	n/a	16	2
	FC-VI-2 port c	1	17	21	n/a	n/a	n/a	17	3
	FC-VI-2 port d	2	17	21	n/a	n/a	n/a	17	3
	HBA port a	1	2	2	2	2	2	2	8
	HBA port b	2	2	2	2	2	2	2	8
	HBA port c	1	3	3	3	3	3	3	9
	HBA port d	2	3	3	3	3	3	3	9

Component	Port	Connect s to FC switch...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller_x_2	FC-VI port a	1	4	4	4	4	4	4	4
	FC-VI port b	2	4	4	4	4	4	4	4
	FC-VI port c	1	5	5	5	5	5	5	5
	FC-VI port d	2	5	5	5	5	5	5	5
	FC-VI-2 port a	1	18	22	n/a	n/a	n/a	20	6
	FC-VI-2 port b	2	18	22	n/a	n/a	n/a	20	6
	FC-VI-2 port c	1	19	23	n/a	n/a	n/a	21	7
	FC-VI-2 port d	2	19	23	n/a	n/a	n/a	21	7
	HBA port a	1	6	6	6	6	6	6	12
	HBA port b	2	6	6	6	6	6	6	12
	HBA port c	1	7	7	7	7	7	7	13
	HBA port d	2	7	7	7	7	7	7	13

MetroCluster 2 or DR group 2

The following table shows the supported controller configurations in MetroCluster 2 or DR group 2 on Brocade switches.

Component	Port	Connect s to FC switch...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller_x_3	FC-VI port a	1	n/a	24	48	n/a	12	18	18
	FC-VI port b	2	n/a	24	48	n/a	12	18	18
	FC-VI port c	1	n/a	25	49	n/a	13	19	19
	FC-VI port d	2	n/a	25	49	n/a	13	19	19
	FC-VI-2 port a	1	n/a	36	n/a	n/a	n/a	36	24
	FC-VI-2 port b	2	n/a	36	n/a	n/a	n/a	36	24
	FC-VI-2 port c	1	n/a	37	n/a	n/a	n/a	37	25
	FC-VI-2 port d	2	n/a	37	n/a	n/a	n/a	37	25
	HBA port a	1	n/a	26	50	n/a	14	24	26
	HBA port b	2	n/a	26	50	n/a	14	24	26
	HBA port c	1	n/a	27	51	n/a	15	25	27
	HBA port d	2	n/a	27	51	n/a	15	25	27

Component	Port	Connect s to FC switch...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
controller_x_4	FC-VI port a	1	n/a	28	52	n/a	16	22	22
	FC-VI port b	2	n/a	28	52	n/a	16	22	22
	FC-VI port c	1	n/a	29	53	n/a	17	23	23
	FC-VI port d	2	n/a	29	53	n/a	17	23	23
	FC-VI-2 port a	1	n/a	38	n/a	n/a	n/a	38	28
	FC-VI-2 port b	2	n/a	38	n/a	n/a	n/a	38	28
	FC-VI-2 port c	1	n/a	39	n/a	n/a	n/a	39	29
	FC-VI-2 port d	2	n/a	39	n/a	n/a	n/a	39	29
	HBA port a	1	n/a	30	54	n/a	18	28	30
	HBA port b	2	n/a	30	54	n/a	18	28	30
	HBA port c	1	n/a	31	55	n/a	19	29	31
	HBA port d	2	n/a	31	55	n/a	19	29	31

MetroCluster 3 or DR group 3

The following table shows the supported controller configurations in MetroCluster 3 or DR group 3 on Brocade switches.

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
controller_x_5	FC-VI port a	1	48	48
	FC-VI port b	2	48	48
	FC-VI port c	1	49	49
	FC-VI port d	2	49	49
	FC-VI-2 port a	1	64	50
	FC-VI-2 port b	2	64	50
	FC-VI-2 port c	1	65	51
	FC-VI-2 port d	2	65	51
	HBA port a	1	50	56
	HBA port b	2	50	56
	HBA port c	1	51	57
	HBA port d	2	51	57

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
controller_x_6	FC-VI port a	1	52	52
	FC-VI port b	2	52	52
	FC-VI port c	1	53	53
	FC-VI port d	2	53	53
	FC-VI-2 port a	1	68	54
	FC-VI-2 port b	2	68	54
	FC-VI-2 port c	1	69	55
	FC-VI-2 port d	2	69	55
	HBA port a	1	54	60
	HBA port b	2	54	60
	HBA port c	1	55	61
	HBA port d	2	55	61

MetroCluster 4 or DR group 4

The following table shows the supported controller configurations in MetroCluster 4 or DR group 4 on Brocade switches.

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
controller_x_7	FC-VI port a	1	66	66
	FC-VI port b	2	66	66
	FC-VI port c	1	67	67
	FC-VI port d	2	67	67
	FC-VI-2 port a	1	84	72
	FC-VI-2 port b	2	84	72
	FC-VI-2 port c	1	85	73
	FC-VI-2 port d	2	85	73
	HBA port a	1	72	74
	HBA port b	2	72	74
	HBA port c	1	73	75
	HBA port d	2	73	75

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
controller_x_8	FC-VI port a	1	70	70
	FC-VI port b	2	70	70
	FC-VI port c	1	71	71
	FC-VI port d	2	71	71
	FC-VI-2 port a	1	86	76
	FC-VI-2 port b	2	86	76
	FC-VI-2 port c	1	87	77
	FC-VI-2 port d	2	87	77
	HBA port a	1	76	78
	HBA port b	2	76	78
	HBA port c	1	77	79
	HBA port d	2	77	79

Brocade port usage for FC-to-SAS bridges in a MetroCluster FC configuration

Learn about the port assignments required to cable Brocade FC switches to FC-to-SAS bridges. The port assignments vary depending on whether the bridges use one or two FC ports.



Brocade 7810 switches only support one DR group.

Shelf configurations using FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2)

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6505, G610, G710, G620, G620-1, G630, and G630-1 switches, you can cable additional bridges to ports 12-15.
- On Brocade 6510 and DCX 8510-8 switches, you can cable additional bridges to ports 12-19.
- On Brocade 6520 switches, you can cable additional bridges to ports 12-21 and 24-45.

- On Brocade 7810 and 7840 switches, MetroCluster 1 or DR group 1 only supports two bridge stacks.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 16-21.

Component		Port	Connec ts to FC switch ...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1	bridge_x _1a	FC1	1	8	8	8	8	8	8	10
		FC2	2	8	8	8	8	8	8	10
	bridge_x _1b	FC1	1	9	9	9	9	9	9	11
		FC2	2	9	9	9	9	9	9	11
Stack 2	bridge_x _2a	FC1	1	10	10	10	10	10	10	14
		FC2	2	10	10	10	10	10	10	14
	bridge_x _2b	FC1	1	11	11	11	11	11	11	15
		FC2	2	11	11	11	11	11	11	15

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6510, DCX 8510-8, and 7840 switches, MetroCluster 2 or DR group 2 only supports two bridge stacks.
- On Brocade 6520 switches, you can cable additional bridges to ports 60-69 and 72-93.
- On Brocade G620, G620-1, G630, and G630-1 switches, you can cable additional bridges to ports 32-35.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 36-39.
- Port usage for the Brocade 6505, G610, and G710 switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, you assign ports on a site-by-site basis depending on the controller model and the number of ISLs and bridge pairs that you're using.

Component		Port	Connects to FC switch ...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1	bridge_x_1a	FC1	1	n/a	32	56	n/a	20	26	32
		FC2	2	n/a	32	56	n/a	20	26	32
	bridge_x_1b	FC1	1	n/a	33	57	n/a	21	27	33
		FC2	2	n/a	33	57	n/a	21	27	33
Stack 2	bridge_x_2a	FC1	1	n/a	34	58	n/a	22	30	34
		FC2	2	n/a	34	58	n/a	22	30	34
	bridge_x_2b	FC1	1	n/a	35	59	n/a	23	31	35
		FC2	2	n/a	35	59	n/a	23	31	35

MetroCluster 3 or DR group 3

The following table shows the supported shelf configurations in MetroCluster 3 or DR group 3 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 60-63.
- On Brocade G730 switches, you can cable additional bridges to ports 64, 65, 68, and 69.

Component		Port	Connects to FC switch...	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	FC1	1	56	58
		FC2	2	56	58
	bridge_x_1b	FC1	1	57	59
		FC2	2	57	59

Component		Port	Connects to FC switch...	G630, G630-1 port	G730 port
Stack 2	bridge_x_2a	FC1	1	58	62
		FC2	2	58	62
	bridge_x_2b	FC1	1	59	63
		FC2	2	59	63

MetroCluster 4 or DR group 4

The following table shows the supported shelf configurations in MetroCluster 4 or DR group 4 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 80-83.
- On Brocade G730 switches, you can cable additional bridges to ports 84-95.

Component		Port	Connects to FC switch...	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	FC1	1	74	80
		FC2	2	74	80
	bridge_x_1b	FC1	1	75	81
		FC2	2	75	81
Stack 2	bridge_x_2a	FC1	1	78	82
		FC2	2	78	82
	bridge_x_2b	FC1	1	79	83
		FC2	2	79	83

Shelf configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 using FibreBridge 7500N or 7600N and only one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6505, G610, G710, G620, G620-1, G630, and G630-1 switches, additional bridges ports 12-

15.

- On Brocade 6510 and DCX 8510-8 switches, you can cable additional bridges to ports 12-19.
- On Brocade 6520 switches, you can cable additional bridges to ports 16-21 and 24-45.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 16-21.

Component	Port	Connects to FC switch...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1	bridge_x_1a	1	8	8	8	8	8	8	10
	bridge_x_1b	2	8	8	8	8	8	8	10
Stack 2	bridge_x_2a	1	9	9	9	9	9	9	11
	bridge_x_2b	2	9	9	9	9	9	9	11
Stack 3	bridge_x_3a	1	10	10	10	10	10	10	14
	bridge_x_3b	2	10	10	10	10	10	10	14
Stack 4	bridge_x_4a	1	11	11	11	11	11	11	15
	bridge_x_4b	2	11	11	11	11	11	11	15

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade 6520 switches, you can cable additional bridges to ports 60-69 and 72-93.
- On Brocade G620, G620-1, G630, G630-1 switches, you can cable additional bridges to ports 32-35.
- On Brocade G720 and G730 switches, you can cable additional bridges to ports 36-39.
- Port usage for the Brocade 6505, G610, and G710 switches in an eight-node MetroCluster configuration is not shown. Due to the limited number of ports, you assign ports on a site-by-site basis depending on the controller model and the number of ISLs and bridge pairs that you're using.

Component	Port	Connects to FC switch...	6505, G610, G710 port	6510, DCX 8510-8 port	6520 port	7810 port	7840 port	G620, G620-1, G630, G630-1 port	G720, G730 port
Stack 1	bridge_x_1a	1	n/a	32	56	n/a	20	26	32
	bridge_x_1b	2	n/a	32	56	n/a	20	26	32
Stack 2	bridge_x_2a	1	n/a	33	57	n/a	21	27	33
	bridge_x_2b	2	n/a	33	57	n/a	21	27	33
Stack 3	bridge_x_3a	1	n/a	34	58	n/a	22	30	34
	bridge_x_3b	2	n/a	34	58	n/a	22	30	34
Stack 4	bridge_x_4a	1	n/a	35	59	n/a	23	31	35
	bridge_x_4b	2	n/a	35	59	n/a	23	31	35

MetroCluster 3 or DR group 3

The following table shows the supported shelf configurations in MetroCluster 3 or DR group 3 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 60-63.
- On Brocade G730 switches, you can cable additional bridges to ports 64, 65, 68, 69.

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	1	56	58
	bridge_x_1b	2	56	58

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
Stack 2	bridge_x_2a	1	57	59
	bridge_x_2b	2	57	59
Stack 3	bridge_x_3a	1	58	62
	bridge_x_3b	2	58	62
Stack 4	bridge_x_4a	1	59	63
	bridge_x_4b	2	59	63

MetroCluster 4 or DR group 4

The following table shows the supported shelf configurations in MetroCluster 4 or DR group 4 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Brocade switches. You should be aware of the following when using this configuration table:

- On Brocade G630 and G630-1 switches, you can cable additional bridges to ports 80-83.
- On Brocade G730 switches, you can cable additional bridges to ports 84-95.

Component	Port	Connects to FC switch...	G630, G630-1 port	G730 port
Stack 1	bridge_x_1a	1	74	80
	bridge_x_1b	2	74	80
Stack 2	bridge_x_2a	1	75	81
	bridge_x_2b	2	75	81
Stack 3	bridge_x_3a	1	78	82
	bridge_x_3b	2	78	82
Stack 4	bridge_x_4a	1	79	83
	bridge_x_4b	2	79	83

Brocade port usage for ISLs in a MetroCluster FC configuration

Learn about the port assignments required to cable Brocade FC switches to ISLs.



- AFF A900 and FAS9500 systems support eight ISLs. Eight ISLs are supported on the Brocade 6510, G620, G620-1, G630, G630-1, G720, and G730 switches.
- Brocade 6520 switches supports eight ISLs, but do not support AFF A900 and FAS9500 systems.

ISL port	6505, G610, G710 port	6520 port	7810 port	7840 (10-Gbps) port	7840 (40-Gbps) port	6510, G620, G620-1, G630, G630-1, G720, G730 port
ISL port 1	20	22	ge2	ge2	ge0	40
ISL port 2	21	23	ge3	ge3	ge1	41
ISL port 3	22	46	ge4	ge10	n/a	42
ISL port 4	23	47	ge5	ge11	n/a	43
ISL port 5	n/a	70	ge6	n/a	n/a	44
ISL port 6	n/a	71	ge7	n/a	n/a	45
ISL port 7	n/a	94	n/a	n/a	n/a	46
ISL port 8	n/a	95	n/a	n/a	n/a	47

Cisco port usage for controllers in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9124V, 9148S, 9148V, 9250i, and 9396S FC switches to your controllers.

The tables show the maximum supported configurations, with eight controller modules in two DR groups. For smaller configurations, ignore the rows for the additional controller modules.



- For Cisco 9132T, see [Cisco 9132T port usage for controllers in a MetroCluster FC configuration](#).
- Cisco 9124V and 9250i switches are not supported in eight-node MetroCluster configurations.

MetroCluster 1 or DR group 1

The following table shows the supported controller configurations in MetroCluster 1 or DR group 1 on Cisco switches (excluding 9132T).

Component	Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x_1	FC-VI port a	1	1	1	1	1	1
	FC-VI port b	2	1	1	1	1	1
	FC-VI port c	1	2	2	2	2	2
	FC-VI port d	2	2	2	2	2	2
	FC-VI-2 port a	1	3	n/a	3	n/a	n/a
	FC-VI-2 port b	2	3	n/a	3	n/a	n/a
	FC-VI-2 port c	1	4	n/a	4	n/a	n/a
	FC-VI-2 port d	2	4	n/a	4	n/a	n/a
	HBA port a	1	13	3	13	3	3
	HBA port b	2	13	3	13	3	3
	HBA port c	1	14	4	14	4	4
	HBA port d	2	14	4	14	4	4

Component	Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x_2	FC-VI port a	1	5	5	5	5	5
	FC-VI port b	2	5	5	5	5	5
	FC-VI port c	1	6	6	6	6	6
	FC-VI port d	2	6	6	6	6	6
	FC-VI-2 port a	1	7	n/a	7	n/a	n/a
	FC-VI-2 port b	2	7	n/a	7	n/a	n/a
	FC-VI-2 port c	1	8	n/a	8	n/a	n/a
	FC-VI-2 port d	2	8	n/a	8	n/a	n/a
	HBA port a	1	15	7	15	7	7
	HBA port b	2	15	7	15	7	7
	HBA port c	1	16	8	16	8	8
	HBA port d	2	16	8	16	8	8

MetroCluster 2 or DR group 2

The following table shows the supported controller configurations in MetroCluster 2 or DR group 2 on Cisco switches (excluding 9132T).

Component	Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x_3	FC-VI port a	1	n/a	25	25	n/a	49
	FC-VI port b	2	n/a	25	25	n/a	49
	FC-VI port c	1	n/a	26	26	n/a	50
	FC-VI port d	2	n/a	26	26	n/a	50
	FC-VI-2 port a	1	n/a	n/a	27	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	27	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	28	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	28	n/a	n/a
	HBA port a	1	n/a	27	37	n/a	51
	HBA port b	2	n/a	27	37	n/a	51
	HBA port c	1	n/a	28	38	n/a	52
	HBA port d	2	n/a	28	38	n/a	52

Component	Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
controller_x_4	FC-VI port a	1	n/a	29	29	n/a	53
	FC-VI port b	2	n/a	29	29	n/a	53
	FC-VI port c	1	n/a	30	30	n/a	54
	FC-VI port d	2	n/a	30	30	n/a	54
	FC-VI-2 port a	1	n/a	n/a	31	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	31	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	32	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	32	n/a	n/a
	HBA port a	1	n/a	31	39	n/a	55
	HBA port b	2	n/a	31	39	n/a	55
	HBA port c	1	n/a	32	40	n/a	56
	HBA port d	1	n/a	32	40	n/a	56

Cisco port usage for FC-to-SAS bridges in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9124V, 9148S, 9148V, 9250i, and 9396S FC switches to FC-to-SAS bridges. The port assignments vary depending on whether the bridges use one or two FC ports.



For Cisco 9132T, see [Cisco 9132t port usage for FC-to-SAS bridges in a MetroCluster FC configuration](#).

Shelf configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 using FibreBridge 7500N or 7600N bridges and both FC ports (FC1 and FC2) on Cisco switches (excluding 9132T). You should be aware of the following when using this configuration table:

- On Cisco 9250i switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-40.
- On Cisco 9396S switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-32.

Component		Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_1a	FC1	1	17	9	17	9	9
		FC2	2	17	9	17	9	9
	bridge_x_1b	FC1	1	18	10	18	10	10
		FC2	2	18	10	18	10	10
Stack 2	bridge_x_2a	FC1	1	19	11	19	11	11
		FC2	2	19	11	19	11	11
	bridge_x_2b	FC1	1	20	12	20	12	12
		FC2	2	20	12	20	12	12
Stack 3	bridge_x_3a	FC1	1	21	13	21	13	13
		FC2	2	21	13	21	13	13
	bridge_x_3b	FC1	1	22	14	22	14	14
		FC2	2	22	14	22	14	14
Stack 4	bridge_x_4a	FC1	1	23	15	23	15	15
		FC2	2	23	15	23	15	15
	bridge_x_4b	FC1	1	24	16	24	16	16
		FC2	2	24	16	24	16	16

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 using FibreBridge 7500N or 7600N and both FC ports (FC1 and FC2) on Cisco switches (excluding 9132T). You should be aware of the following when using the cabling tables:

- Cisco 9124V and 9250i switches are not supported for eight-node MetroCluster configurations.

- On Cisco 9396S switches, you can cable additional MetroCluster 2 (DR group 2) bridges to ports 65-80.

Component		Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_1a	FC1	1	n/a	33	41	n/a	57
		FC2	2	n/a	33	41	n/a	57
	bridge_x_1b	FC1	1	n/a	34	42	n/a	58
		FC2	2	n/a	34	42	n/a	58
Stack 2	bridge_x_2a	FC1	1	n/a	35	43	n/a	59
		FC2	2	n/a	35	43	n/a	59
	bridge_x_2b	FC1	1	n/a	36	44	n/a	60
		FC2	2	n/a	36	44	n/a	60
Stack 3	bridge_x_3a	FC1	1	n/a	37	45	n/a	61
		FC2	2	n/a	37	45	n/a	61
	bridge_x_3b	FC1	1	n/a	38	46	n/a	62
		FC2	2	n/a	38	46	n/a	62
Stack 4	bridge_x_4a	FC1	1	n/a	39	47	n/a	63
		FC2	2	n/a	39	47	n/a	63
	bridge_x_4b	FC1	1	n/a	40	48	n/a	64
		FC2	2	n/a	40	48	n/a	64

Shelf configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Cisco switches (excluding 9132T). The reference configuration file (RCF) doesn't support one FC port on FibreBridge bridges, so you must configure the back-end fibre channel switches manually.

[Configure the Cisco FC switches manually](#)

You should be aware of the following when using the cabling tables:

- On Cisco 9250i switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-40.
- On Cisco 9396S switches, you can cable additional MetroCluster 1 or DR group 1 bridges to ports 17-32.

Component	Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_1a	1	17	9	17	9	9
	bridge_x_1b	2	17	9	17	9	9
Stack 2	bridge_x_2a	1	18	10	18	10	10
	bridge_x_2b	2	18	10	18	10	10
Stack 3	bridge_x_3a	1	19	11	19	11	11
	bridge_x_3b	2	19	11	19	11	11
Stack 4	bridge_x_4a	1	20	12	20	12	12
	bridge_x_4b	2	20	12	20	12	12
Stack 5	bridge_x_5a	1	21	13	21	13	13
	bridge_x_5b	2	21	13	21	13	13
Stack 6	bridge_x_6a	1	22	14	22	14	14
	bridge_x_6b	2	22	14	22	14	14
Stack 7	bridge_x_7a	1	23	15	23	15	15
	bridge_x_7b	2	23	15	23	15	15
Stack 8	bridge_x_8a	1	24	16	24	16	16
	bridge_x_8b	2	24	16	24	16	16

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using one FC port (FC1 or FC2) on Cisco switches (excluding 9132T). You should be aware of the following when using this configuration table:

- The Cisco 9124V and 9250i switches are not supported for eight-node MetroCluster configurations.
- On Cisco 9396S switches, you can cable additional MetroCluster 2 or DR group 2 bridges to ports 65-80.

Component	Port	Connects to FC switch...	9124V port	9148S port	9148V port	9250i port	9396S port
Stack 1	bridge_x_1a	1	n/a	33	41	n/a	57
	bridge_x_1b	2	n/a	33	41	n/a	57
Stack 2	bridge_x_2a	1	n/a	34	42	n/a	58
	bridge_x_2b	2	n/a	34	42	n/a	58
Stack 3	bridge_x_3a	1	n/a	35	43	n/a	59
	bridge_x_3b	2	n/a	35	43	n/a	59
Stack 4	bridge_x_4a	1	n/a	36	44	n/a	60
	bridge_x_4b	2	n/a	36	44	n/a	60
Stack 5	bridge_x_5a	1	n/a	37	45	n/a	61
	bridge_x_5b	2	n/a	37	45	n/a	61
Stack 6	bridge_x_6a	1	n/a	38	46	n/a	62
	bridge_x_6b	2	n/a	38	46	n/a	62
Stack 7	bridge_x_7a	1	n/a	39	47	n/a	63
	bridge_x_7b	2	n/a	39	47	n/a	63
Stack 8	bridge_x_8a	1	n/a	40	48	n/a	64
	bridge_x_8b	2	n/a	40	48	n/a	64

Cisco port usage for ISLs in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9124V, 9148S, 9148V, 9250i, and 9396S FC switches to ISLs.

The following table shows ISL port usage. ISL port usage is the same on all switches in the configuration.



- For Cisco 9132T, see [Cisco 9132T port usage for ISLs in a MetroCluster FC configuration](#).
- The Cisco 9250i switch requires a 24 port license.

ISL port	9124V port	9148S port	9148V port	9250i port	9396S port
ISL port 1	9	20	9	12	44
ISL port 2	10	24	10	16	48
ISL port 3	11	44	11	20	92
ISL port 4	12	48	12	24	96
ISL port 5	n/a	n/a	33	n/a	n/a
ISL port 6	n/a	n/a	34	n/a	n/a
ISL port 7	n/a	n/a	35	n/a	n/a
ISL port 8	n/a	n/a	36	n/a	n/a

Cisco 9132T port usage for controllers in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9132T FC switches to your controllers.

The following table shows controller configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2). The tables show the maximum supported configurations with four and eight controller modules in two DR groups.



For eight-node configurations, you must perform the zoning manually because RCFs are not provided.

MetroCluster 1 or DR group 1

The following table shows the supported controller configurations for MetroCluster 1 or DR group 1 on Cisco 9132T switches. You should be aware of the following when using this configuration table:

- AFF A900 and FAS9500 systems have eight FC-VI ports (a, b, c, and d for FC-VI-1 and FC-VI-2).

Component	Port	Connects to FC_switch...	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
controller_x_1	FC-VI port a	1	LEM1-1	LEM1-1	LEM1-1
	FC-VI port b	2	LEM1-1	LEM1-1	LEM1-1
	FC-VI port c	1	LEM1-2	LEM1-2	LEM1-2
	FC-VI port d	2	LEM1-2	LEM1-2	LEM1-2
	FC-VI-2 port a	1	LEM1-3	LEM1-3	n/a
	FC-VI-2 port b	2	LEM1-3	LEM1-3	n/a
	FC-VI-2 port c	1	LEM1-4	LEM1-4	n/a
	FC-VI-2 port d	2	LEM1-4	LEM1-4	n/a
	HBA port a	1	LEM1-5	LEM1-5	LEM1-3
	HBA port b	2	LEM1-5	LEM1-5	LEM1-3
	HBA port c	1	LEM1-6	LEM1-6	LEM1-4
	HBA port d	2	LEM1-6	LEM1-6	LEM1-4

Component	Port	Connects to FC_switch...	9132T 1x LEM (Four-node)	9132T 2x LEM (Four-node)	9132T 2x LEM (Eight-node)
controller_x_2	FC-VI port a	1	LEM1-7	LEM1-7	LEM1-5
	FC-VI port b	2	LEM1-7	LEM1-7	LEM1-5
	FC-VI port c	1	LEM1-8	LEM1-8	LEM1-6
	FC-VI port d	2	LEM1-8	LEM1-8	LEM1-6
	FC-VI-2 port a	1	LEM1-9	LEM1-9	n/a
	FC-VI-2 port b	2	LEM1-9	LEM1-9	n/a
	FC-VI-2 port c	1	LEM1-10	LEM1-10	n/a
	FC-VI-2 port d	2	LEM1-10	LEM1-10	n/a
	HBA port a	1	LEM1-11	LEM1-11	LEM1-7
	HBA port b	2	LEM1-11	LEM1-11	LEM1-7
	HBA port c	1	LEM1-12	LEM1-12	LEM1-8
	HBA port d	2	LEM1-12	LEM1-12	LEM1-8

MetroCluster 2 or DR group 2

The following table shows the supported Cisco 9132T controller configurations for MetroCluster 2 or DR group 2 on Cisco 9132T switches. You should be aware of the following when using this configuration table:

- AFF A900 and FAS9500 systems have eight FC-VI ports (a, b, c, and d for FC-VI-1 and FC-VI-2).
- MetroCluster 2 or DR group 2 is not supported on Cisco 9132T switches for AFF A900 and FAS9500 systems.
- MetroCluster 2 or DR group 2 is only supported in eight-node MetroCluster configurations

Component	Port	Connects to FC_switch...	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
controller_x_3	FC-VI port a	1	n/a	n/a	LEM2-1
	FC-VI port b	2	n/a	n/a	LEM2-1
	FC-VI port c	1	n/a	n/a	LEM2-2
	FC-VI port d	2	n/a	n/a	LEM2-2
	FC-VI-2 port a	1	n/a	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	n/a
	HBA port a	1	n/a	n/a	LEM2-3
	HBA port b	2	n/a	n/a	LEM2-3
	HBA port c	1	n/a	n/a	LEM2-4
	HBA port d	2	n/a	n/a	LEM2-4

Component	Port	Connects to FC_switch...	9132T 1x LEM (Four-node)	9132T 2x LEM (Four-node)	9132T 2x LEM (Eight-node)
controller_x_4	FC-VI-1 port a	1	n/a	n/a	LEM2-5
	FC-VI-1 port b	2	n/a	n/a	LEM2-5
	FC-VI-1 port c	1	n/a	n/a	LEM2-6
	FC-VI-1 port d	2	n/a	n/a	LEM2-6
	FC-VI-2 port a	1	n/a	n/a	n/a
	FC-VI-2 port b	2	n/a	n/a	n/a
	FC-VI-2 port c	1	n/a	n/a	n/a
	FC-VI-2 port d	2	n/a	n/a	n/a
	HBA port a	1	n/a	n/a	LEM2-7
	HBA port b	2	n/a	n/a	LEM2-7
	HBA port c	1	n/a	n/a	LEM2-8
	HBA port d	2	n/a	n/a	LEM2-8

Cisco 9132T port usage for FC-to-SAS bridges in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9132T FC switches to FC-to-SAS bridges using both FC ports.



Only one (1) bridge stack is supported using Cisco 9132T switches with 1xLEM Module.

MetroCluster 1 or DR group 1

The following table shows the supported shelf configurations in MetroCluster 1 or DR group 1 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Cisco 9132T switches. You should be aware of the following when using this configuration table:

- In four-node configurations, you can cable additional bridges to ports LEM2-1 through LEM2-8 on Cisco 9132T switches with 2xLEMs.

Component		Port	Connects to FC_switch...	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
Stack 1	bridge_x_1a	FC1	1	LEM1-13	LEM1-13	LEM1-9
		FC2	2	LEM1-13	LEM1-13	LEM1-9
	bridge_x_1b	FC1	1	LEM1-14	LEM1-14	LEM1-10
		FC2	2	LEM1-14	LEM1-14	LEM1-10
Stack 2	bridge_x_2a	FC1	1	n/a	LEM1-15	LEM1-11
		FC2	2	n/a	LEM1-15	LEM1-11
	bridge_x_2b	FC1	1	n/a	LEM1-16	LEM1-12
		FC2	2	n/a	LEM1-16	LEM1-12

MetroCluster 2 or DR group 2

The following table shows the supported shelf configurations in MetroCluster 2 or DR group 2 for FibreBridge 7500N or 7600N bridges using both FC ports (FC1 and FC2) on Cisco 9132T switches. You should be aware of the following when using this configuration table:

- In eight-node configurations, you can cable additional bridges to ports LEM2-13 through LEM2-16 on Cisco 9132T switches with 2x LEMs.

Component		Port	Connects to FC_switch...	9132T 1x LEM (Four- node)	9132T 2x LEM (Four- node)	9132T 2x LEM (Eight- node)
Stack 1	bridge_x_1a	FC1	1	n/a	n/a	LEM1-9
		FC2	2	n/a	n/a	LEM1-9
	bridge_x_1b	FC1	1	n/a	n/a	LEM1-10
		FC2	2	n/a	n/a	LEM1-10

Component		Port	Connects to FC_switch...	9132T 1x LEM (Four-node)	9132T 2x LEM (Four-node)	9132T 2x LEM (Eight-node)
Stack 2	bridge_x_2a	FC1	1	n/a	n/a	LEM1-11
		FC2	2	n/a	n/a	LEM1-11
	bridge_x_2b	FC1	1	n/a	n/a	LEM1-12
		FC2	2	n/a	n/a	LEM1-12

Cisco 9132T port usage for ISLs in a MetroCluster FC configuration

Learn about the port assignments required to cable Cisco 9132T FC switches to ISLs.

The following table shows ISL port usage for a Cisco 9132T switch.

ISL port	9132T 1x LEM (Four-node)	9132T 2x LEM (Four-node)	9132T 2x LEM (Eight-node)
ISL port 1	LEM1-15	LEM2-9	LEM1-13
ISL port 2	LEM1-16	LEM2-10	LEM1-14
ISL port 3	n/a	LEM2-11	LEM1-15
ISL port 4	n/a	LEM2-12	LEM1-16
ISL port 5	n/a	LEM2-13	n/a
ISL port 6	n/a	LEM2-14	n/a
ISL port 7	n/a	LEM2-15	n/a
ISL port 8	n/a	LEM2-16	n/a

Cabling the cluster interconnect in eight- or four-node configurations

In eight-node or four-node MetroCluster configurations, you must cable the cluster interconnect between the local controller modules at each site.

About this task

This task is not required on two-node MetroCluster configurations.

This task must be performed at both MetroCluster sites.

Step

1. Cable the cluster interconnect from one controller module to the other, or if cluster interconnect switches are used, from each controller module to the switches.

Related information

[ONTAP Hardware Systems Documentation](#)

[Network and LIF management](#)

Cabling the cluster peering connections

You must cable the controller module ports used for cluster peering so that they have connectivity with the cluster on the partner site.

About this task

This task must be performed on each controller module in the MetroCluster configuration.

At least two ports on each controller module should be used for cluster peering.

The recommended minimum bandwidth for the ports and network connectivity is 1 GbE.

Step

1. Identify and cable at least two ports for cluster peering and verify they have network connectivity with the partner cluster.

Cluster peering can be done on dedicated ports or on data ports. Using dedicated ports provides higher throughput for the cluster peering traffic.

Related information

[Cluster and SVM peering express configuration](#)

Each MetroCluster site is configured as a peer to its partner site. You should be familiar with the prerequisites and guidelines for configuring the peering relationships and when deciding whether to use shared or dedicated ports for those relationships.

[Cluster peering](#)

Cabling the HA interconnect

If you have an eight- or a four-node MetroCluster configuration and the storage controllers within the HA pairs are in separate chassis, you must cable the HA interconnect between the controllers.

About this task

- This task does not apply to two-node MetroCluster configurations.
- This task must be performed at both MetroCluster sites.
- The HA interconnect must be cabled only if the storage controllers within the HA pair are in separate chassis.

Some storage controller models support two controllers in a single chassis, in which case they use an internal HA interconnect.

Steps

1. Cable the HA interconnect if the storage controller's HA partner is in a separate chassis.

[ONTAP Hardware Systems Documentation](#)

2. If the MetroCluster site includes two HA pairs, repeat the previous steps on the second HA pair.
3. Repeat this task at the MetroCluster partner site.

Cabling the management and data connections

You must cable the management and data ports on each storage controller to the site networks.

About this task

This task must be repeated for each new controller at both MetroCluster sites.

You can connect the controller and cluster switch management ports to existing switches in your network or to new dedicated network switches such as NetApp CN1601 cluster management switches.

Step

1. Cable the controller's management and data ports to the management and data networks at the local site.

[ONTAP Hardware Systems Documentation](#)

Configure the FC switches

FC switch configuration overview

You can configure Cisco and Brocade FC switches by using RCF files, or, if necessary, you can manually configure the switches.

If you...	Use the procedure...
Have an RCF that meets your requirements	<ul style="list-style-type: none">• Configure Brocade FC switches with RCF files• Configure Cisco FC switches with RCF files
Do not have an RCF or have an RCF that does not meet your requirements	<ul style="list-style-type: none">• Configure the Brocade FC switches manually• Configure the Cisco FC switches manually

Configure Brocade FC switches with RCF files

Resetting the Brocade FC switch to factory defaults

Before installing a new software version and RCF files, you must erase the current switch configuration and perform basic configuration.

About this task

You must repeat these steps on each of the FC switches in the MetroCluster fabric configuration.

Steps

1. Log in to the switch as an administrator.
2. Disable the Brocade Virtual Fabrics (VF) feature:

```
fosconfig options
```

```
FC_switch_A_1:admin> fosconfig --disable vf
WARNING: This is a disruptive operation that requires a reboot to take
effect.
Would you like to continue [Y/N]: y
```

3. Disconnect the ISL cables from the ports on the switch.
4. Disable the switch:

```
switchcfgpersistentdisable
```

```
FC_switch_A_1:admin> switchcfgpersistentdisable
```

5. Disable the configuration:

```
cfgDisable
```

```
FC_switch_A_1:admin> cfgDisable
You are about to disable zoning configuration. This action will disable
any previous zoning configuration enabled.
Do you want to disable zoning configuration? (yes, y, no, n): [no] y
Updating flash ...
Effective configuration is empty. "No Access" default zone mode is ON.
```

6. Clear the configuration:

```
cfgClear
```

```
FC_switch_A_1:admin> cfgClear
The Clear All action will clear all Aliases, Zones, FA Zones
and configurations in the Defined configuration.
Run cfgSave to commit the transaction or cfgTransAbort to
cancel the transaction.
Do you really want to clear all configurations? (yes, y, no, n): [no] y
```

7. Save the configuration:

```
cfgSave
```

```
FC_switch_A_1:admin> cfgSave
You are about to save the Defined zoning configuration. This
action will only save the changes on Defined configuration.
Do you want to save the Defined zoning configuration only? (yes, y, no,
n): [no] y
Updating flash ...
```

8. Set the default configuration:

```
configDefault
```

```
FC_switch_A_1:admin> configDefault
WARNING: This is a disruptive operation that requires a switch reboot.
Would you like to continue [Y/N]: y
Executing configdefault...Please wait
2020/10/05-08:04:08, [FCR-1069], 1016, FID 128, INFO, FC_switch_A_1, The
FC Routing service is enabled.
2020/10/05-08:04:08, [FCR-1068], 1017, FID 128, INFO, FC_switch_A_1, The
FC Routing service is disabled.
2020/10/05-08:04:08, [FCR-1070], 1018, FID 128, INFO, FC_switch_A_1, The
FC Routing configuration is set to default.
Committing configuration ... done.
2020/10/05-08:04:12, [MAPS-1113], 1019, FID 128, INFO, FC_switch_A_1,
Policy dflt_conservative_policy activated.
2020/10/05-08:04:12, [MAPS-1145], 1020, FID 128, INFO, FC_switch_A_1,
FPI Profile dflt_fpi_profile is activated for E-Ports.
2020/10/05-08:04:12, [MAPS-1144], 1021, FID 128, INFO, FC_switch_A_1,
FPI Profile dflt_fpi_profile is activated for F-Ports.
The switch has to be rebooted to allow the changes to take effect.
2020/10/05-08:04:12, [CONF-1031], 1022, FID 128, INFO, FC_switch_A_1,
configDefault completed successfully for switch.
```

9. Set the port configuration to default for all ports:

```
portcfgdefault port-number
```

```
FC_switch_A_1:admin> portcfgdefault <port number>
```

You must complete this step for each port.

10. If you are running a version earlier than FOS 9.0, verify that the switch is using the dynamic Port on Demand (POD) method.



In Fabric OS 9.0 and later, the license method is dynamic by default. The static license method is not supported.

For Brocade Fabric OS versions before 8.0, you run the following commands as admin, and for versions 8.0 and later, you run them as root.

- a. Run the license command:

```
licenseport --show.
```

```
FC_switch_A_1:admin> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
```

- b. Enable the root user if it is disabled by Brocade.

```
FC_switch_A_1:admin> userconfig --change root -e yes
FC_switch_A_1:admin> rootaccess --set consoleonly
```

- c. Run the license command:

```
licenseport --show.
```

```
FC_switch_A_1:root> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
```

- d. If you are running Fabric OS 8.2.x and earlier, you must change the license method to dynamic:

```
licenseport --method dynamic
```

```
FC_switch_A_1:admin> licenseport --method dynamic
The POD method has been changed to dynamic.
Please reboot the switch now for this change to take effect
```

11. Reboot the switch:

```
fastBoot
```

```
FC_switch_A_1:admin> fastboot
Warning: This command would cause the switch to reboot
and result in traffic disruption.
Are you sure you want to reboot the switch [y/n]?y
```

12. Confirm that the default settings have been implemented:

```
switchShow
```

13. Verify that the IP address is set correctly:

```
ipAddrShow
```

You can set the IP address with the following command, if required:

```
ipAddrSet
```

Downloading the Brocade FC switch RCF file

You must download the reference configuration (RCF) file to each switch in the MetroCluster fabric configuration.

About this task

To use these RCF files, the system must be running ONTAP 9.1 or later and you must use the port layout for ONTAP 9.1 or later.

If you are planning to use only one of the FC ports on the FibreBridge bridges, configure the back-end fibre channel switches manually using the instructions found in the section, [Port assignments for FC switches](#).

Steps

1. Refer to the RCF file table on the Brocade RCF download page and identify the correct RCF file for each switch in your configuration.

The RCF files must be applied to the correct switches.

2. Download the RCF files for the switches from the [MetroCluster RCF download](#) page.

The files must be placed in a location where they can be transferred to the switch. There is a separate file for each of the four switches that make up the two-switch fabric.

3. Repeat these steps on each switch in the configuration.

Installing the Brocade FC switch RCF file

When you configure a Brocade FC switch, you can install the switch configuration files that provide the complete switch settings for certain configurations.

About this task

- You must repeat these steps on each of the Brocade FC switches in the MetroCluster fabric configuration.
- If you use an xWDM configuration, you might require additional settings on the ISLs. See the xWDM

vendor documentation for more information.

Steps

1. Initiate the download and configuration process:

```
configDownload
```

Respond to the prompts as shown in the following example.

```
FC_switch_A_1:admin> configDownload
Protocol (scp, ftp, sftp, local) [ftp]:
Server Name or IP Address [host]: <user input>
User Name [user]:<user input>
Path/Filename [<home dir>/config.txt]:path to configuration file
Section (all|chassis|switch [all]): all
.
.
.
Do you want to continue [y/n]: y
Password: <user input>
```

After entering your password, the switch downloads and executes the configuration file.

2. Confirm that the configuration file has set the switch domain:

```
switchShow
```

Each switch is assigned a different domain number depending on which configuration file the switch used.

```
FC_switch_A_1:admin> switchShow
switchName: FC_switch_A_1
switchType: 109.1
switchState: Online
switchMode: Native
switchRole: Subordinate
switchDomain: 5
```

3. Verify that your switch is assigned the correct domain value as indicated in the following table.

Fabric	Switch	Switch domain
1	A_1	5
	B_1	7

2	A_2	6
	B_2	8

4. Change the port speed:

`portcfgspeed`

```
FC_switch_A_1:admin> portcfgspeed port number port speed
```

By default, all the ports are configured to operate at 16 Gbps. You might change the port speed for the following reasons:

- The interconnect switch ports speed should be changed when an 8-Gbps FC-VI adapter is used and the switch port speed should set to 8 Gbps.
- The ISL ports' speed must be changed when the ISL is not capable of running at 16 Gbps.

5. Calculate the ISL distance.

Due to the behavior of the FC-VI, you must set the distance to 1.5 times the real distance with a minimum of 10 (LE). The distance for the ISL is calculated as follows, rounded up to the next full kilometer: $1.5 \times \text{real distance} = \text{distance}$.

If the distance is 3 km, then $1.5 \times 3 \text{ km} = 4.5$. This is lower than 10; therefore, you must set the ISL to the LE distance level.

The distance is 20 km, then $1.5 \times 20 \text{ km} = 30$. You must set the ISL to the LS distance level.

6. Set the distance for each ISL port:

`portcfglongdistance port level vc_link_init -distance distance_value`

A `vc_link_init` value of 1 uses the fillword "ARB" by default. A value of 0 uses the fillword "IDLE". The required value might vary depending on the link you use. In this example, the default is set and the distance is assumed to be 20 km. Hence, the setting is "30" with a `vc_link_init` value of "1", and the ISL port is "21".

Example: LS

```
FC_switch_A_1:admin> portcfglongdistance 21 LS 1 -distance 30
```

Example: LE

```
FC_switch_A_1:admin> portcfglongdistance 21 LE 1
```

7. Persistently enable the switch:

`switchcfgpersistentenable`

The example shows how to persistently enable FC switch_A_1.

```
FC_switch_A_1:admin> switchcfgpersistentenable
```

8. Verify if the IP address is set correctly:

```
ipAddrshow
```

```
FC_switch_A_1:admin> ipAddrshow
```

You can set the IP address, if required:

```
ipAddrSet
```

9. Set the timezone from the switch prompt:

```
tstimezone --interactive
```

You should respond to the prompts as required.

```
FC_switch_A_1:admin> tstimezone --interactive
```

10. Reboot the switch:

```
reboot
```

The example shows how to reboot FC switch_A_1.

```
FC_switch_A_1:admin> reboot
```

11. Verify the distance setting:

```
portbuffershow
```

A distance setting of LE appears as 10 km.

```

FC_Switch_A_1:admin> portbuffershow
User Port Lx    Max/Resv Buffer Needed Link      Remaining
Port Type Mode Buffers  Usage  Buffers Distance Buffers
-----
...
21    E    -     8     67    67     30 km
22    E    -     8     67    67     30 km
...
23    -    8     0     -     -     466

```

12. Reconnect the ISL cables to the ports on the switches where they were removed.

The ISL cables were disconnected when the factory settings were reset to the default settings.

Resetting the Brocade FC switch to factory defaults

13. Validate the configuration.

a. Verify that the switches form one fabric:

```
switchshow
```

The following example shows the output for a configuration that uses ISLs on ports 20 and 21.

```

FC_switch_A_1:admin> switchshow
switchName: FC_switch_A_1
switchType: 109.1
switchState:Online
switchMode: Native
switchRole: Subordinate
switchDomain:      5
switchId:   fffc01
switchWwn:  10:00:00:05:33:86:89:cb
zoning:      OFF
switchBeacon: OFF

Index Port Address Media Speed State Proto
=====
...
20    20  010C00   id    16G  Online FC   LE E-Port
10:00:00:05:33:8c:2e:9a "FC_switch_B_1" (downstream) (trunk master)
21    21  010D00   id    16G  Online FC   LE E-Port (Trunk port,
master is Port 20)
...

```

b. Confirm the configuration of the fabrics:

fabricshow

```
FC_switch_A_1:admin> fabricshow
  Switch ID      Worldwide Name          Enet IP Addr FC IP Addr Name
-----
1: fffc01 10:00:00:05:33:86:89:cb 10.10.10.55  0.0.0.0
"FC_switch_A_1"
3: fffc03 10:00:00:05:33:8c:2e:9a 10.10.10.65  0.0.0.0
>"FC_switch_B_1"
```

c. Verify that the ISLs are working:

islshow

```
FC_switch_A_1:admin> islshow
```

d. Confirm that zoning is properly replicated:

cfgshow

zonestow

Both outputs should show the same configuration information and zoning information for both switches.

e. If trunking is used, confirm the trunking:

trunkShow

```
FC_switch_A_1:admin> trunkshow
```

Configure the Cisco FC switches with RCF files

Resetting the Cisco FC switch to factory defaults

Before installing a new software version and RCFs, you must erase the Cisco switch configuration and perform basic configuration.

About this task

You must repeat these steps on each of the FC switches in the MetroCluster fabric configuration.



The outputs shown are for Cisco IP switches; however, these steps are also applicable for Cisco FC switches.

Steps

1. Reset the switch to factory defaults:

a. Erase the existing configuration:

```
write erase
```

b. Reload the switch software:

```
reload
```

The system reboots and enters the configuration wizard. During the boot, if you receive the prompt Abort Auto Provisioning and continue with normal setup?(yes/no)[n], you should respond **yes** to proceed.

c. In the configuration wizard, enter the basic switch settings:

- Admin password
- Switch name
- Out-of-band management configuration
- Default gateway
- SSH service (Remote Support Agent).

After completing the configuration wizard, the switch reboots.

d. When prompted, enter the user name and password to log in to the switch.

The following example shows the prompts and system responses when logging in to the switch. The angle brackets (<<<) show where you enter the information.

```
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:y
**<<<<**

    Enter the password for "admin": password **<<<<**
    Confirm the password for "admin": password **<<<<**
        ---- Basic System Configuration Dialog VDC: 1 ----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Please register Cisco Nexus3000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus3000 devices must be registered to receive
entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.
```

e. Enter basic information in the next set of prompts, including the switch name, management address,

The following configuration will be applied:

```
password strength-check
switchname IP_switch_A_1
vrf context management
ip route 0.0.0.0/0 10.10.99.1
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address 10.10.99.10 255.255.255.0
no shutdown
```

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]:

```
2017 Jun 13 21:24:43 A1 %$ VDC-1 %$ %COPP-2-COPP_POLICY: Control-
Plane is protected with policy copp-system-p-policy-strict.
```

```
[#####] 100%
Copy complete.
```

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
.
.
.
IP_switch_A_1#
```

2. Save the configuration:

```
IP_switch_A_1# copy running-config startup-config
```

3. Reboot the switch and wait for the switch to reload:

```
IP_switch_A_1# reload
```

4. Repeat the previous steps on the other three switches in the MetroCluster fabric configuration.

Downloading and installing the Cisco FC switch NX-OS software

You must download the switch operating system file and RCF file to each switch in the MetroCluster fabric configuration.

Before you begin

This task requires file transfer software, such as FTP, TFTP, SFTP, or SCP, to copy the files to the switches.

About this task

These steps must be repeated on each of the FC switches in the MetroCluster fabric configuration.

You must use the supported switch software version.

NetApp Hardware Universe



The outputs shown are for Cisco IP switches; however, these steps are also applicable for Cisco FC switches.

Steps

1. Download the supported NX-OS software file.

[Cisco download page](#)

2. Copy the switch software to the switch:

```
copy sftp://root@server-ip-address/tftpboot/NX-OS-file-name bootflash: vrf
management
```

In this example, the `nxos.7.0.3.I4.6.bin` file is copied from SFTP server 10.10.99.99 to the local bootflash:

```
IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/nxos.7.0.3.I4.6.bin
bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/nxos.7.0.3.I4.6.bin
/bootflash/nxos.7.0.3.I4.6.bin
Fetching /tftpboot/nxos.7.0.3.I4.6.bin to /bootflash/nxos.7.0.3.I4.6.bin
/tftpboot/nxos.7.0.3.I4.6.bin          100% 666MB 7.2MB/s
01:32
sftp> exit
Copy complete, now saving to disk (please wait)...
```

3. Verify on each switch that the switch NX-OS files are present in each switch's bootflash directory:

```
dir bootflash
```

The following example shows that the files are present on `IP_switch_A_1`:

```
IP_switch_A_1# dir bootflash:
      .
      .
      .
698629632   Jun 13 21:37:44 2017  nxos.7.0.3.I4.6.bin
      .
      .
      .

Usage for bootflash://sup-local
 1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#
```

4. Install the switch software:

```
install all system bootflash:nxos.version-number.bin kickstart
bootflash:nxos.version-kickstart-number.bin
```

```
IP_switch_A_1# install all system bootflash:nxos.7.0.3.I4.6.bin
kickstart bootflash:nxos.7.0.3.I4.6.bin
Installer will perform compatibility check first. Please wait.

Verifying image bootflash:/nxos.7.0.3.I4.6.bin for boot variable
"kickstart".
[#####] 100% -- SUCCESS

Verifying image bootflash:/nxos.7.0.3.I4.6.bin for boot variable
"system".
[#####] 100% -- SUCCESS

Performing module support checks.
[#####] 100% -- SUCCESS

Verifying image type.
[#####] 100% -- SUCCESS

Extracting "system" version from image bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS

Extracting "kickstart" version from image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS
...
```

The switch reboot automatically after the switch software has installed.

5. Wait for the switch to reload and then log in to the switch.

After the switch has rebooted the login prompt is displayed:

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
.
MDP database restore in progress.
IP_switch_A_1#

The switch software is now installed.
```

6. Verify that the switch software has been installed:

```
show version
```

The following example shows the output:

```

IP_switch_A_1# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
.

Software
  BIOS: version 04.24
  NXOS: version 7.0(3)I4(6)   **<<< switch software version**
  BIOS compile time: 04/21/2016
  NXOS image file is: bootflash:///nxos.7.0.3.I4.6.bin
  NXOS compile time: 3/9/2017 22:00:00 [03/10/2017 07:05:18]

Hardware
  cisco Nexus 3132QV Chassis
  Intel(R) Core(TM) i3- CPU @ 2.50GHz with 16401416 kB of memory.
  Processor Board ID FOC20123GPS

  Device name: A1
  bootflash: 14900224 kB
  usb1: 0 kB (expansion flash)

Kernel uptime is 0 day(s), 0 hour(s), 1 minute(s), 49 second(s)

Last reset at 403451 usecs after Mon Jun 10 21:43:52 2017

Reason: Reset due to upgrade
System version: 7.0(3)I4(1)
Service:

plugin
  Core Plugin, Ethernet Plugin
IP_switch_A_1#

```

7. Repeat these steps on the remaining three FC switches in the MetroCluster fabric configuration.

Downloading and installing the Cisco FC RCF files

You must download the RCF file to each switch in the MetroCluster fabric configuration.

Before you begin

This task requires file transfer software, such as FTP, Trivial File Transfer Protocol (TFTP), SFTP, or Secure

Copy Protocol (SCP), to copy the files to the switches.

About this task

These steps must be repeated on each of the Cisco FC switches in the MetroCluster fabric configuration.

You must use the supported switch software version.

[NetApp Hardware Universe](#)

There are four RCF files, one for each of the four switches in the MetroCluster fabric configuration. You must use the correct RCF files for the switch model you are using.

Switch	RCF file
FC_switch_A_1	NX3232_v1.80_Switch-A1.txt
FC_switch_A_2	NX3232_v1.80_Switch-A2.txt
FC_switch_B_1	NX3232_v1.80_Switch-B1.txt
FC_switch_B_2	NX3232_v1.80_Switch-B2.txt



The outputs shown are for Cisco IP switches; however, these steps are also applicable for Cisco FC switches.

Steps

1. Download the Cisco FC RCF files from the [MetroCluster RCF download page](#).
2. Copy the RCF files to the switches.
 - a. Copy the RCF files to the first switch:

```
copy sftp://root@FTP-server-IP-address/tftpboot/switch-specific-RCF
bootflash: vrf management
```

In this example, the `NX3232_v1.80_Switch-A1.txt` RCF file is copied from the SFTP server at `10.10.99.99` to the local bootflash. You must use the IP address of your TFTP/SFTP server and the file name of the RCF file that you need to install.

```

IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/NX3232_v1.8T-
X1_Switch-A1.txt bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/NX3232_v1.80_Switch-A1.txt
/bootflash/NX3232_v1.80_Switch-A1.txt
Fetching /tftpboot/NX3232_v1.80_Switch-A1.txt to
/bootflash/NX3232_v1.80_Switch-A1.txt
/tftpboot/NX3232_v1.80_Switch-A1.txt          100% 5141      5.0KB/s
00:00
sftp> exit
Copy complete, now saving to disk (please wait)...
IP_switch_A_1#

```

- b. Repeat the previous substep for each of the other three switches, being sure to copy the matching RCF file to the corresponding switch.
3. Verify on each switch that the RCF file is present in each switch's `bootflash` directory:

```
dir bootflash:
```

The following example shows that the files are present on `IP_switch_A_1`:

```

IP_switch_A_1# dir bootflash:
      .
      .
      .
5514   Jun 13 22:09:05 2017  NX3232_v1.80_Switch-A1.txt
      .
      .
      .

Usage for bootflash://sup-local
 1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#

```

4. Copy the matching RCF file from the local bootflash to the running configuration on each switch:

```
copy bootflash:switch-specific-RCF.txt running-config
```

5. Copy the RCF files from the running configuration to the startup configuration on each switch:

```
copy running-config startup-config
```

You should see output similar to the following:

```
IP_switch_A_1# copy bootflash:NX3232_v1.80_Switch-A1.txt running-config
IP_switch_A_1# copy running-config startup-config
```

6. Reload the switch:

```
reload
```

```
IP_switch_A_1# reload
```

7. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Configure the Brocade FC switches manually

You must configure each of the Brocade switch fabrics in the MetroCluster configuration.

Before you begin

- You must have a PC or UNIX workstation with Telnet or Secure Shell (SSH) access to the FC switches.
- You must be using four supported Brocade switches of the same model with the same Brocade Fabric Operating System (FOS) version and licensing.

[NetApp Interoperability Matrix Tool](#)

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- The four supported Brocade switches must be connected to two fabrics of two switches each, with each fabric spanning both sites.
- Each storage controller must have four initiator ports available to connect to the switch fabrics. Two initiator ports must be connected from each storage controller to each fabric.



You can configure FAS8020, AFF8020, FAS8200, and AFF A300 systems with two initiators ports per controller (a single initiator port to each fabric) if all the following criteria are met:

- There are fewer than four FC initiator ports available to connect the disk storage and no additional ports can be configured as FC initiators.
- All slots are in use and no FC initiator card can be added.

About this task

- You should enable Inter-Switch Link (ISL) trunking when it is supported by the links.

[Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations](#)

- If you use an xWDM configuration, you might require additional settings on the ISLs. See the xWDM vendor documentation for more information.
- All ISLs must have the same length and same speed in one fabric.

Different lengths can be used in the different fabrics. The same speed must be used in all fabrics.

- Metro-E and TDM (SONET/SDH) are not supported, and any non-FC native framing or signaling is not supported.

Metro-E means Ethernet framing or signaling occurs either natively over a Metro distance or through some time-division multiplexing (TDM), multiprotocol label switching (MPLS), or wavelength-division multiplexing (WDM).

- TDMs, FCR (native FC Routing), or FCIP extensions are not supported for the MetroCluster FC switch fabric.
- The following Brocade FC switches support WAN ISL encryption and compression for MetroCluster FC back-end fabrics:
 - Brocade G720
 - Brocade G630
 - Brocade G620
 - Brocade 6520
- The Brocade Virtual Fabric (VF) feature is not supported.
- FC zoning based on domain port is supported, but zoning based on worldwide name (WWN) is not supported.

Review Brocade license requirements

You need certain licenses for the switches in a MetroCluster configuration. You must install these licenses on all four switches.

About this task

The MetroCluster configuration has the following Brocade license requirements:

- Trunking license for systems using more than one ISL, as recommended.
- Extended Fabric license (for ISL distances over 6 km)
- Enterprise license for sites with more than one ISL and an ISL distance greater than 6 km

The Enterprise license includes Brocade Network Advisor and all licenses except for additional port licenses.

Step

1. Verify that the licenses are installed:

For Fabric OS 8.2.x and earlier

Run the command `licenseshow`.

For Fabric OS 9.0 and later

Run the command `license --show`.

If you do not have these licenses, you should contact your sales representative before proceeding.

Set the Brocade FC switch values to factory defaults

You must set the switch to its factory defaults to ensure a successful configuration. You must also assign each switch a unique name.

About this task

In the examples in this procedure, the fabric consists of BrocadeSwitchA and BrocadeSwitchB.

Steps

1. Make a console connection and log in to both switches in one fabric.
2. Disable the switch persistently:

```
switchcfgpersistentdisable
```

This ensures the switch will remain disabled after a reboot or fastboot. If this command is not available, use the `switchdisable` command.

The following example shows the command on BrocadeSwitchA:

```
BrocadeSwitchA:admin> switchcfgpersistentdisable
```

The following example shows the command on BrocadeSwitchB:

```
BrocadeSwitchB:admin> switchcfgpersistentdisable
```

3. Set the switch name:

```
switchname switch_name
```

The switches should each have a unique name. After setting the name, the prompt changes accordingly.

The following example shows the command on BrocadeSwitchA:

```
BrocadeSwitchA:admin> switchname "FC_switch_A_1"  
FC_switch_A_1:admin>
```

The following example shows the command on BrocadeSwitchB:

```
BrocadeSwitchB:admin> switchname "FC_Switch_B_1"  
FC_switch_B_1:admin>
```

4. Set all ports to their default values:

```
portcfgdefault
```

This must be done for all ports on the switch.

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:admin> portcfgdefault 0
FC_switch_A_1:admin> portcfgdefault 1
...
FC_switch_A_1:admin> portcfgdefault 39
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1:admin> portcfgdefault 0
FC_switch_B_1:admin> portcfgdefault 1
...
FC_switch_B_1:admin> portcfgdefault 39
```

5. Clear the zoning information:

```
cfgdisable
```

```
cfgclear
```

```
cfgsave
```

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:admin> cfgdisable
FC_switch_A_1:admin> cfgclear
FC_switch_A_1:admin> cfgsave
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1:admin> cfgdisable
FC_switch_B_1:admin> cfgclear
FC_switch_B_1:admin> cfgsave
```

6. Set the general switch settings to default:

```
configdefault
```

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> configdefault
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> configdefault
```

7. Set all ports to non-trunking mode:

```
switchcfgtrunk 0
```

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> switchcfgtrunk 0
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> switchcfgtrunk 0
```

8. On Brocade 6510 switches, disable the Brocade Virtual Fabrics (VF) feature:

```
fosconfig options
```

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> fosconfig --disable vf
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> fosconfig --disable vf
```

9. Clear the Administrative Domain (AD) configuration:

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1:> defzone --noaccess  
FC_switch_A_1:> cfgsave  
FC_switch_A_1:> exit
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_A_1:> defzone --noaccess  
FC_switch_A_1:> cfgsave  
FC_switch_A_1:> exit
```

10. Reboot the switch:

reboot

The following example shows the command on FC_switch_A_1:

```
FC_switch_A_1:admin> reboot
```

The following example shows the command on FC_switch_B_1:

```
FC_switch_B_1:admin> reboot
```

Configure basic switch settings

You must configure basic global settings, including the domain ID, for Brocade switches.

About this task

This task contains steps that must be performed on each switch at both of the MetroCluster sites.

In this procedure, you set the unique domain ID for each switch as shown in the following example. In the example, domain IDs 5 and 7 form fabric_1, and domain IDs 6 and 8 form fabric_2.

- FC_switch_A_1 is assigned to domain ID 5
- FC_switch_A_2 is assigned to domain ID 6
- FC_switch_B_1 is assigned to domain ID 7
- FC_switch_B_2 is assigned to domain ID 8

Steps

1. Enter configuration mode:

```
configure
```

2. Proceed through the prompts:

- a. Set the domain ID for the switch.
- b. Press **Enter** in response to the prompts until you get to "RDP Polling Cycle", and then set that value to 0 to disable the polling.
- c. Press **Enter** until you return to the switch prompt.

```

FC_switch_A_1:admin> configure
Fabric parameters = y
Domain_id = 5
.
.

RSCN Transmission Mode [yes, y, no, no: [no] y

End-device RSCN Transmission Mode
(0 = RSCN with single PID, 1 = RSCN with multiple PIDs, 2 = Fabric
RSCN): (0..2) [1]
Domain RSCN To End-device for switch IP address or name change
(0 = disabled, 1 = enabled): (0..1) [0] 1

.
.
RDP Polling Cycle(hours) [0 = Disable Polling]: (0..24) [1] 0

```

3. If you are using two or more ISLs per fabric, then you can configure either in-order delivery (IOD) of frames or out-of-order (OOD) delivery of frames.



The standard IOD settings are recommended. You should configure OOD only if necessary.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

- a. The following steps must be performed on each switch fabric to configure IOD of frames:

- i. Enable IOD:

```
iodset
```

- ii. Set the Advanced Performance Tuning (APT) policy to 1:

```
aptpolicy 1
```

- iii. Disable Dynamic Load Sharing (DLS):

```
dlsreset
```

- iv. Verify the IOD settings by using the `iodshow`, `aptpolicy`, and `dls` commands.

For example, issue the following commands on `FC_switch_A_1`:

```
FC_switch_A_1:admin> iodshow
IOD is set

FC_switch_A_1:admin> aptpolicy
Current Policy: 1 0(ap)

3 0(ap) : Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
    0: AP Shared Link Policy
    1: AP Dedicated Link Policy
command aptpolicy completed

FC_switch_A_1:admin> dlsshow
DLS is not set
```

- v. Repeat these steps on the second switch fabric.
- b. The following steps must be performed on each switch fabric to configure OOD of frames:
 - i. Enable OOD:

```
iodreset
```

- ii. Set the Advanced Performance Tuning (APT) policy to 3:

```
aptopolicy 3
```

- iii. Disable Dynamic Load Sharing (DLS):

```
dlsreset
```

- iv. Verify the OOD settings:

```
iodshow
```

```
aptopolicy
```

```
dlsshow
```

For example, issue the following commands on FC_switch_A_1:

```

FC_switch_A_1:admin> iodshow
IOD is not set

FC_switch_A_1:admin> aptpolicy
Current Policy: 3 0(ap)
3 0(ap) : Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
0: AP Shared Link Policy
1: AP Dedicated Link Policy
command aptpolicy completed

FC_switch_A_1:admin> dlsshow
DLS is set by default with current routing policy

```

- v. Repeat these steps on the second switch fabric.



When configuring ONTAP on the controller modules, OOD must be explicitly configured on each controller module in the MetroCluster configuration.

Configure in-order delivery or out-of-order delivery of frames on ONTAP software

- 4. If you are running a version earlier than FOS 9.0, verify that the switch is using the dynamic Port on Demand (POD) licensing method.



In Fabric OS 9.0 and later, the license method is dynamic by default. The static license method is not supported.

- a. Run the license command:

```
licenseport --show
```

```

FC_switch_A_1:admin> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use

```



Brocade FabricOS versions before 8.0 run the following commands as admin and versions 8.0 and later run them as root.

- b. Enable the root user.

If the root user is already disabled by Brocade, enable the root user as shown in the following example:

```
FC_switch_A_1:admin> userconfig --change root -e yes
FC_switch_A_1:admin> rootaccess --set consoleonly
```

c. Run the license command:

```
license --show -port
```

```
FC_switch_A_1:root> license --show -port
24 ports are available in this switch
Full POD license is installed
Dynamic POD method is in use
```

d. If you are running Fabric OS 8.2.x and earlier, you must change the license method to dynamic:

```
licenseport --method dynamic
```

```
FC_switch_A_1:admin> licenseport --method dynamic
The POD method has been changed to dynamic.
Please reboot the switch now for this change to take effect
```

5. Enable the trap for T11-FC-ZONE-SERVER-MIB to provide successful health monitoring of the switches in ONTAP:

a. Enable the T11-FC-ZONE-SERVER-MIB:

```
snmpconfig --set mibCapability -mib_name T11-FC-ZONE-SERVER-MIB -bitmask
0x3f
```

b. Enable the T11-FC-ZONE-SERVER-MIB trap:

```
snmpconfig --enable mibcapability -mib_name SW-MIB -trap_name
swZoneConfigChangeTrap
```

c. Repeat the previous steps on the second switch fabric.

6. **Optional:** If you set the community string to a value other than "public", you must configure the ONTAP Health Monitors using the community string you specify:

a. Change the existing community string:

```
snmpconfig --set snmpv1
```

b. Press **Enter** until you see "Community (ro): [public]" text.

c. Enter the desired community string.

On FC_switch_A_1:

```

FC_switch_A_1:admin> snmpconfig --set snmpv1
SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [private]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [public] mcchm      <<<<<< change the community string
to the desired value,
Trap Recipient's IP address : [0.0.0.0]      in this example it is set
to "mcchm"
Community (ro): [common]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address : [0.0.0.0]
Committing configuration.....done.
FC_switch_A_1:admin>

```

On FC_switch_B_1:

```

FC_switch_B_1:admin> snmpconfig --set snmpv1
SNMP community and trap recipient configuration:
Community (rw): [Secret C0de]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [OrigEquipMfr]
Trap Recipient's IP address : [0.0.0.0]
Community (rw): [private]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [public] mcchm      <<<<<< change the community
string to the desired value,
Trap Recipient's IP address : [0.0.0.0]      in this example it is set
to "mcchm"
Community (ro): [common]
Trap Recipient's IP address : [0.0.0.0]
Community (ro): [FibreChannel]
Trap Recipient's IP address : [0.0.0.0]
Committing configuration.....done.
FC_switch_B_1:admin>

```

7. Reboot the switch:

```
reboot
```

On FC_switch_A_1:

```
FC_switch_A_1:admin> reboot
```

On FC_switch_B_1:

```
FC_switch_B_1:admin> reboot
```

8. Persistently enable the switch:

```
switchcfgpersistentenable
```

On FC_switch_A_1:

```
FC_switch_A_1:admin> switchcfgpersistentenable
```

On FC_switch_B_1:

```
FC_switch_B_1:admin> switchcfgpersistentenable
```

Configure basic switch settings on a Brocade DCX 8510-8 switch

You must configure basic global settings, including the domain ID, for Brocade switches.

About this task

You must perform the steps on each switch at both MetroCluster sites. In this procedure, you set the domain ID for each switch as shown in the following examples:

- FC_switch_A_1 is assigned to domain ID 5
- FC_switch_A_2 is assigned to domain ID 6
- FC_switch_B_1 is assigned to domain ID 7
- FC_switch_B_2 is assigned to domain ID 8

In the previous example, domain IDs 5 and 7 form fabric_1, and domain IDs 6 and 8 form fabric_2.



You can also use this procedure to configure the switches when you are only using one DCX 8510-8 switch per site.

Using this procedure, you should create two logical switches on each Brocade DCX 8510-8 switch. The two logical switches created on both Brocade DCX8510-8 switches will form two logical fabrics as shown in the following examples:

- LOGICAL FABRIC 1: Switch1/Blade1 and Switch 2 Blade 1
- LOGICAL FABRIC 2: Switch1/Blade2 and Switch 2 Blade 2

Steps

1. Enter the command mode:

```
configure
```

2. Proceed through the prompts:

- a. Set the domain ID for the switch.
- b. Keep selecting **Enter** until you get to "RDP Polling Cycle", and then set the value to 0 to disable the polling.
- c. Select **Enter** until you return to the switch prompt.

```
FC_switch_A_1:admin> configure
Fabric parameters = y
Domain_id = `5

RDP Polling Cycle(hours) [0 = Disable Polling]: (0..24) [1] 0
`
```

3. Repeat these steps on all switches in fabric_1 and fabric_2.

4. Configure the virtual fabrics.

- a. Enable virtual fabrics on the switch:

```
fosconfig --enablevf
```

- b. Configure the system to use the same base configuration on all logical switches:

```
configurechassis
```

The following example shows the output for the `configurechassis` command:

```
System (yes, y, no, n): [no] n
cfgload attributes (yes, y, no, n): [no] n
Custom attributes (yes, y, no, n): [no] y
Config Index (0 to ignore): (0..1000) [3]:
```

5. Create and configure the logical switch:

```
scfg --create fabricID
```

6. Add all ports from a blade to the virtual fabric:

```
lscfg --config fabricID -slot slot -port lowest-port - highest-port
```



The blades forming a logical fabric (e.g. Switch 1 Blade 1 and Switch 3 Blade 1) need to have the same fabric ID.

```

setcontext fabricid
switchdisable
configure
<configure the switch per the above settings>
switchname unique switch name
switchenable

```

Related information

[Requirements for using a Brocade DCX 8510-8 switch](#)

Configure E-ports on Brocade FC switches using FC ports

For Brocade switches on which the Inter-Switch Links (ISL) are configured using FC ports, you must configure the switch ports on each switch fabric that connect the ISL. These ISL ports are also known as E-ports.

Before you begin

- All of the ISLs in an FC switch fabric must be configured with the same speed and distance.
- The combination of the switch port and small form-factor pluggable (SFP) must support the speed.
- The supported ISL distance depends on the FC switch model.

[NetApp Interoperability Matrix Tool](#)

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- The ISL link must have a dedicated lambda, and the link must be supported by Brocade for the distance, switch type, and Fabric Operating System (FOS).

About this task

You must not use the L0 setting when issuing the `portCfgLongDistance` command. Instead, you should use the LE or LS setting to configure the distance on the Brocade switches with a minimum of LE distance level.

You must not use the LD setting when issuing the `portCfgLongDistance` command when working with xWDM/TDM equipment. Instead, you should use the LE or LS setting to configure the distance on the Brocade switches.

You must perform this task for each FC switch fabric.

The following tables show the ISL ports for different switches and different number of ISLs in a configuration running ONTAP 9.1 or 9.2. The examples shown in this section are for a Brocade 6505 switch. You should modify the examples to use ports that apply to your switch type.

You must use the required number of ISLs for your configuration.

Switch model	ISL port	Switch port
--------------	----------	-------------

Brocade 6520	ISL port 1	23
	ISL port 2	47
	ISL port 3	71
	ISL port 4	95
Brocade 6505	ISL port 1	20
	ISL port 2	21
	ISL port 3	22
	ISL port 4	23
Brocade 6510 and Brocade DCX 8510-8	ISL port 1	40
	ISL port 2	41
	ISL port 3	42
	ISL port 4	43
	ISL port 5	44
	ISL port 6	45
	ISL port 7	46
	ISL port 8	47
Brocade 7810	ISL port 1	ge2 (10-Gbps)
	ISL port 2	ge3(10-Gbps)
	ISL port 3	ge4 (10-Gbps)
	ISL port 4	ge5 (10-Gbps)
	ISL port 5	ge6 (10-Gbps)
	ISL port 6	ge7 (10-Gbps)
Brocade 7840	ISL port 1	ge0 (40-Gbps) or ge2 (10-Gbps)
	ISL port 2	ge1 (40-Gbps) or ge3 (10-Gbps)
	ISL port 3	ge10 (10-Gbps)
	ISL port 4	ge11 (10-Gbps)

Note: The Brocade 7840 switch supports either two 40 Gbps VE-ports or up to four 10 Gbps VE-ports per switch for the creation of FCIP ISLs.

Brocade G610, G710	ISL port 1	20
	ISL port 2	21
	ISL port 3	22
	ISL port 4	23
Brocade G620, G620-1, G630, G630-1, G720	ISL port 1	40
	ISL port 2	41
	ISL port 3	42
	ISL port 4	43
	ISL port 5	44
	ISL port 6	45
	ISL port 7	46

Steps

1. Configure the port speed:

```
portcfgspeed port-numberspeed
```

You must use the highest common speed that is supported by the components in the path.

In the following example, there are two ISLs for each fabric:

```
FC_switch_A_1:admin> portcfgspeed 20 16
FC_switch_A_1:admin> portcfgspeed 21 16

FC_switch_B_1:admin> portcfgspeed 20 16
FC_switch_B_1:admin> portcfgspeed 21 16
```

2. Configure the trunking mode for each ISL:

```
portcfgtrunkport port-number
```

- If you are configuring the ISLs for trunking (IOD), set the portcfgtrunk port-numberport-number to 1 as shown in the following example:

```

FC_switch_A_1:admin> portcfgtrunkport 20 1
FC_switch_A_1:admin> portcfgtrunkport 21 1
FC_switch_B_1:admin> portcfgtrunkport 20 1
FC_switch_B_1:admin> portcfgtrunkport 21 1

```

- If you do not want to configure the ISL for trunking (OOD), set portcfgtrunkport-number to 0 as shown in the following example:

```

FC_switch_A_1:admin> portcfgtrunkport 20 0
FC_switch_A_1:admin> portcfgtrunkport 21 0
FC_switch_B_1:admin> portcfgtrunkport 20 0
FC_switch_B_1:admin> portcfgtrunkport 21 0

```

3. Enable QoS traffic for each of the ISL ports:

```
portcfgqos --enable port-number
```

In the following example, there are two ISLs per switch fabric:

```

FC_switch_A_1:admin> portcfgqos --enable 20
FC_switch_A_1:admin> portcfgqos --enable 21

FC_switch_B_1:admin> portcfgqos --enable 20
FC_switch_B_1:admin> portcfgqos --enable 21

```

4. Verify the settings:

```
portCfgShow command
```

The following example shows the output for a configuration that uses two ISLs cabled to port 20 and port 21. The Trunk Port setting should be ON for IOD and OFF for OOD:

```

Ports of Slot 0  12  13  14 15  16  17  18  19  20  21 22  23  24
25  26  27
-----+---+---+---+---+---+---+---+---+---+---+---+---+
-----+---+---+---
Speed          AN  AN  AN  AN  AN  AN  8G  AN  AN  AN  16G  16G
AN  AN  AN  AN
Fill Word      0  0  0  0  0  0  3  0  0  0  3  3  3
0  0  0
AL_PA Offset 13 ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..  ..
..  ..  ..  ..
Trunk Port     ..  ..  ..  ..  ..  ..  ..  ..  ON  ON  ..  ..
..  ..  ..  ..

```

Long Distance
..													
VC Link Init
..													
Locked L_Port
..													
Locked G_Port
..													
Disabled E_Port
..													
Locked E_Port
..													
ISL R_RDY Mode
..													
RSCN Suppressed
..													
Persistent Disable
..													
LOS TOV enable
..													
NPIV capability	ON												
ON ON ON ON													
NPIV PP Limit	126	126	126	126	126	126	126	126	126	126	126	126	126
126 126 126 126													
QOS E_Port	AE												
AE AE AE AE													
Mirror Port
..													
Rate Limit
..													
Credit Recovery	ON												
ON ON ON ON													
Fport Buffers
..													
Port Auto Disable
..													
CSCTL mode
..													
Fault Delay	0	0	0	0	0	0	0	0	0	0	0	0	0

5. Calculate the ISL distance.

Because of the behavior of FC-VI, the distance must be set to 1.5 times the real distance with a minimum distance of 10 km (using the LE distance level).

The distance for the ISL is calculated as follows, rounded up to the next full kilometer:

$1.5 \times \text{real_distance} = \text{distance}$

If the distance is 3 km, then $1.5 \times 3 \text{ km} = 4.5 \text{ km}$. This is lower than 10 km, so the ISL must be set to the LE distance level.

If the distance is 20 km, then $1.5 \times 20 \text{ km} = 30 \text{ km}$. The ISL must be set to 30 km and must use the LS distance level.

6. Set the distance on each ISL port:

```
portcfglongdistance portdistance-level vc_link_init distance
```

A `vc_link_init` value of 1 uses the ARB fill word (default). A value of 0 uses IDLE. The required value might depend on the link being used. The commands must be repeated for each ISL port.

For an ISL distance of 3 km, as given in the example in the previous step, the setting is 4.5 km with the default `vc_link_init` value of 1. Because a setting of 4.5 km is lower than 10 km, the port needs to be set to the LE distance level:

```
FC_switch_A_1:admin> portcfglongdistance 20 LE 1  
  
FC_switch_B_1:admin> portcfglongdistance 20 LE 1
```

For an ISL distance of 20 km, as given in the example in the previous step, the setting is 30 km with the default `vc_link_init` value of 1:

```
FC_switch_A_1:admin> portcfglongdistance 20 LS 1 -distance 30  
  
FC_switch_B_1:admin> portcfglongdistance 20 LS 1 -distance 30
```

7. Verify the distance setting:

```
portbuffershow
```

A distance level of LE appears as 10 km.

The following example shows the output for a configuration that uses ISLs on port 20 and port 21:

```
FC_switch_A_1:admin> portbuffershow
```

User Port	Port Type	Lx Mode	Max/Resv Buffers	Buffer Usage	Needed Buffers	Link Distance	Remaining Buffers
----	----	----	-----	-----	-----	-----	-----
...							
20	E	-	8	67	67	30km	
21	E	-	8	67	67	30km	
...							
23		-	8	0	-	-	466

8. Verify that both switches form one fabric:

```
switchshow
```

The following example shows the output for a configuration that uses ISLs on port 20 and port 21:

```

FC_switch_A_1:admin> switchshow
switchName: FC_switch_A_1
switchType: 109.1
switchState:Online
switchMode: Native
switchRole: Subordinate
switchDomain:      5
switchId:   fffc01
switchWwn:  10:00:00:05:33:86:89:cb
zoning:     OFF
switchBeacon: OFF

Index Port Address Media Speed State Proto
=====
...
20  20  010C00  id    16G  Online FC  LE E-Port
10:00:00:05:33:8c:2e:9a "FC_switch_B_1" (downstream) (trunk master)
21  21  010D00  id    16G  Online FC  LE E-Port (Trunk port, master
is Port 20)
...

FC_switch_B_1:admin> switchshow
switchName: FC_switch_B_1
switchType: 109.1
switchState:Online
switchMode: Native
switchRole: Principal
switchDomain:      7
switchId:   fffc03
switchWwn:  10:00:00:05:33:8c:2e:9a
zoning:     OFF
switchBeacon: OFF

Index Port Address Media Speed State Proto
=====
...
20  20  030C00  id    16G  Online FC  LE E-Port
10:00:00:05:33:86:89:cb "FC_switch_A_1" (downstream) (Trunk master)
21  21  030D00  id    16G  Online FC  LE E-Port (Trunk port, master
is Port 20)
...

```

9. Confirm the configuration of the fabrics:

```
fabricshow
```

```

FC_switch_A_1:admin> fabricshow
  Switch ID      Worldwide Name      Enet IP Addr FC IP Addr Name
-----
1: fffc01 10:00:00:05:33:86:89:cb 10.10.10.55  0.0.0.0
"FC_switch_A_1"
3: fffc03 10:00:00:05:33:8c:2e:9a 10.10.10.65  0.0.0.0
>"FC_switch_B_1"

```

```

FC_switch_B_1:admin> fabricshow
  Switch ID      Worldwide Name      Enet IP Addr FC IP Addr      Name
-----
1: fffc01 10:00:00:05:33:86:89:cb 10.10.10.55  0.0.0.0
"FC_switch_A_1"

3: fffc03 10:00:00:05:33:8c:2e:9a 10.10.10.65  0.0.0.0
>"FC_switch_B_1"

```

10. Confirm the trunking of the ISLs:

trunkshow

- If you are configuring the ISLs for trunking (IOD), you should see output similar to the following:

```

FC_switch_A_1:admin> trunkshow
  1: 20-> 20 10:00:00:05:33:ac:2b:13 3 deskew 15 MASTER
    21-> 21 10:00:00:05:33:8c:2e:9a 3 deskew 16
FC_switch_B_1:admin> trunkshow
  1: 20-> 20 10:00:00:05:33:86:89:cb 3 deskew 15 MASTER
    21-> 21 10:00:00:05:33:86:89:cb 3 deskew 16

```

- If you are not configuring the ISLs for trunking (OOD), you should see output similar to the following:

```

FC_switch_A_1:admin> trunkshow
  1: 20-> 20 10:00:00:05:33:ac:2b:13 3 deskew 15 MASTER
  2: 21-> 21 10:00:00:05:33:8c:2e:9a 3 deskew 16 MASTER
FC_switch_B_1:admin> trunkshow
  1: 20-> 20 10:00:00:05:33:86:89:cb 3 deskew 15 MASTER
  2: 21-> 21 10:00:00:05:33:86:89:cb 3 deskew 16 MASTER

```

11. Repeat [Step 1](#) through [Step 10](#) for the second FC switch fabric.

Related information

[Port assignments for FC switches](#)

Configuring 10 Gbps VE ports on Brocade FC 7840 switches

When using the 10 Gbps VE ports (which use FCIP) for ISLs, you must create IP interfaces on each port, and configure FCIP tunnels and circuits in each tunnel.

About this task

This procedure must be performed on each switch fabric in the MetroCluster configuration.

The examples in this procedure assume that the two Brocade 7840 switches have the following IP addresses:

- FC_switch_A_1 is local.
- FC_switch_B_1 is remote.

Steps

1. Create IP interface (ipif) addresses for the 10 Gbps ports on both switches in the fabric:

```
portcfg ipif FC_switch1_namefirst_port_name create FC_switch1_IP_address
netmask netmask_number vlan 2 mtu auto
```

The following command creates ipif addresses on ports ge2.dp0 and ge3.dp0 of FC_switch_A_1:

```
portcfg ipif ge2.dp0 create 10.10.20.71 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge3.dp0 create 10.10.21.71 netmask 255.255.0.0 vlan 2 mtu
auto
```

The following command creates ipif addresses on ports ge2.dp0 and ge3.dp0 of FC_switch_B_1:

```
portcfg ipif ge2.dp0 create 10.10.20.72 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge3.dp0 create 10.10.21.72 netmask 255.255.0.0 vlan 2 mtu
auto
```

2. Verify that the ipif addresses were created successfully on both switches:

```
portshow ipif all
```

The following command shows the ipif addresses on switch FC_switch_A_1:

```
FC_switch_A_1:root> portshow ipif all
```

Port	IP Address	/ Pfx	MTU	VLAN	Flags
ge2.dp0	10.10.20.71	/ 24	AUTO	2	U R M I
ge3.dp0	10.10.21.71	/ 20	AUTO	2	U R M I

Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

The following command shows the ipif addresses on switch FC_switch_B_1:

```
FC_switch_B_1:root> portshow ipif all
```

Port	IP Address	/ Pfx	MTU	VLAN	Flags
ge2.dp0	10.10.20.72	/ 24	AUTO	2	U R M I
ge3.dp0	10.10.21.72	/ 20	AUTO	2	U R M I

Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport

3. Create the first of the two FCIP tunnels using the ports on dp0:

```
portcfg fciptunnel
```

This command creates a tunnel with a single circuit.

The following command creates the tunnel on switch FC_switch_A_1:

```
portcfg fciptunnel 24 create -S 10.10.20.71 -D 10.10.20.72 -b 10000000  
-B 10000000
```

The following command creates the tunnel on switch FC_switch_B_1:

```
portcfg fciptunnel 24 create -S 10.10.20.72 -D 10.10.20.71 -b 10000000  
-B 10000000
```

4. Verify that the FCIP tunnels were successfully created:

```
portshow fcip tunnel all
```

The following example shows that the tunnels were created and the circuits are up:

```
FC_switch_B_1:root>

 Tunnel Circuit  OpStatus  Flags      Uptime    TxMBps    RxMBps    ConnCnt
CommRt Met/G
-----
-----
 24    -          Up        -----   2d8m     0.05     0.41     3        -
-----
-----
Flags (tunnel): i=IPSec f=Fastwrite T=TapePipelining F=FICON
r=ReservedBW
                  a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                  I=IP-Ext
```

5. Create an additional circuit for dp0.

The following command creates a circuit on switch FC_switch_A_1 for dp0:

```
portcfg fcipcircuit 24 create 1 -S 10.10.21.71 -D 10.10.21.72 --min
-comm-rate 5000000 --max-comm-rate 5000000
```

The following command creates a circuit on switch FC_switch_B_1 for dp0:

```
portcfg fcipcircuit 24 create 1 -S 10.10.21.72 -D 10.10.21.71 --min
-comm-rate 5000000 --max-comm-rate 5000000
```

6. Verify that all circuits were successfully created:

```
portshow fcipcircuit all
```

The following command shows the circuits and their status:

```
FC_switch_A_1:root> portshow fcipcircuit all
```

Tunnel CommRt	Circuit Met/G	OpStatus	Flags	Uptime	TxMBps	RxMBps	ConnCnt
24 10000/10000	0 ge2 0/-	Up	---va---4	2d12m	0.02	0.03	3
24 10000/10000	1 ge3 0/-	Up	---va---4	2d12m	0.02	0.04	3

Flags (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4
6=IPv6
ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA

Configure 40 Gbps VE-ports on Brocade 7810 and 7840 FC switches

When using the two 40 GbE VE-ports (which use FCIP) for ISLs, you must create IP interfaces on each port, and configure FCIP tunnels and circuits in each tunnel.

About this task

This procedure must be performed on each switch fabric in the MetroCluster configuration.

The examples in this procedure use two switches:

- FC_switch_A_1 is local.
- FC_switch_B_1 is remote.

Steps

1. Create IP interface (ipif) addresses for the 40 Gbps ports on both switches in the fabric:

```
portcfg ipif FC_switch_namefirst_port_name create FC_switch_IP_address netmask  
netmask_number vlan 2 mtu auto
```

The following command creates ipif addresses on ports ge0.dp0 and ge1.dp0 of FC_switch_A_1:

```
portcfg ipif ge0.dp0 create 10.10.82.10 netmask 255.255.0.0 vlan 2 mtu  
auto  
portcfg ipif ge1.dp0 create 10.10.82.11 netmask 255.255.0.0 vlan 2 mtu  
auto
```

The following command creates ipif addresses on ports ge0.dp0 and ge1.dp0 of FC_switch_B_1:

```
portcfg ipif ge0.dp0 create 10.10.83.10 netmask 255.255.0.0 vlan 2 mtu
auto
portcfg ipif ge1.dp0 create 10.10.83.11 netmask 255.255.0.0 vlan 2 mtu
auto
```

2. Verify that the ipif addresses were successfully created on both switches:

```
portshow ipif all
```

The following example shows the IP interfaces on FC_switch_A_1:

```
Port          IP Address          / Pfx  MTU   VLAN  Flags
-----
---
-----
ge0.dp0       10.10.82.10         / 16   AUTO  2     U R M
ge1.dp0       10.10.82.11         / 16   AUTO  2     U R M
-----
-----
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
      N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

The following example shows the IP interfaces on FC_switch_B_1:

```
Port          IP Address          / Pfx  MTU   VLAN  Flags
-----
-----
ge0.dp0       10.10.83.10         / 16   AUTO  2     U R M
ge1.dp0       10.10.83.11         / 16   AUTO  2     U R M
-----
-----
Flags: U=Up B=Broadcast D=Debug L=Loopback P=Point2Point R=Running
I=InUse
      N=NoArp PR=Promisc M=Multicast S=StaticArp LU=LinkUp X=Crossport
```

3. Create the FCIP tunnel on both switches:

```
portcfg fciptunnel
```

The following command creates the tunnel on FC_switch_A_1:

```
portcfg fciptunnel 24 create -S 10.10.82.10 -D 10.10.83.10 -b 10000000
-B 10000000
```

The following command creates the tunnel on FC_switch_B_1:

```
portcfg fciptunnel 24 create -S 10.10.83.10 -D 10.10.82.10 -b 10000000
-B 10000000
```

4. Verify that the FCIP tunnel has been successfully created:

```
portshow fciptunnel all
```

The following example shows that the tunnel was created and the circuits are up:

```
FC_switch_A_1:root>

 Tunnel Circuit  OpStatus  Flags      Uptime    TxMBps   RxMBps  ConnCnt
CommRt Met/G
-----
-----
 24      -          Up        -----   2d8m     0.05    0.41   3       -
-
-----
-----
Flags (tunnel): i=IPSec f=Fastwrite T=TapePipelining F=FICON
r=ReservedBW
                  a=FastDeflate d=Deflate D=AggrDeflate P=Protocol
                  I=IP-Ext
```

5. Create an additional circuit on each switch:

```
portcfg fcipcircuit 24 create 1 -S source-IP-address -D destination-IP-address
--min-comm-rate 10000000 --max-comm-rate 10000000
```

The following command creates a circuit on switch FC_switch_A_1 for dp0:

```
portcfg fcipcircuit 24 create 1 -S 10.10.82.11 -D 10.10.83.11 --min
-comm-rate 10000000 --max-comm-rate 10000000
```

The following command creates a circuit on switch FC_switch_B_1 for dp1:

```
portcfg fcipcircuit 24 create 1 -S 10.10.83.11 -D 10.10.82.11 --min
-comm-rate 10000000 --max-comm-rate 10000000
```

6. Verify that all circuits were successfully created:

```
portshow fcipcircuit all
```

The following example lists the circuits and shows that their OpStatus is up:

```
FC_switch_A_1:root> portshow fcipcircuit all

 Tunnel Circuit  OpStatus  Flags      Uptime  TxMBps  RxMBps  ConnCnt
CommRt  Met/G
-----
-----
 24    0 ge0      Up        ---va---4  2d12m   0.02    0.03    3
10000/10000 0/-
 24    1 ge1      Up        ---va---4  2d12m   0.02    0.04    3
10000/10000 0/-
-----
-----
Flags (circuit): h=HA-Configured v=VLAN-Tagged p=PMTU i=IPSec 4=IPv4
6=IPv6
                ARL a=Auto r=Reset s=StepDown t=TimedStepDown S=SLA
```

Configure the non-E-ports on the Brocade switch

You must configure the non-E-ports on the FC switch. In a MetroCluster configuration, these are the ports that connect the switch to the HBA initiators, FC-VI interconnects, and FC-to-SAS bridges. These steps must be done for each port.

About this task

In the following example, the ports connect an FC-to-SAS bridge:

- Port 6 on FC_FC_switch_A_1 at Site_A
- Port 6 on FC_FC_switch_B_1 at Site_B

Steps

1. Configure the port speed for each non-E-port:

```
portcfgspeed portspeed
```

You should use the highest common speed, which is the highest speed supported by all components in the data path: the SFP, the switch port that the SFP is installed on, and the connected device (HBA, bridge, and so on).

For example, the components might have the following supported speeds:

- The SFP is capable of 4, 8, or 16 GB.
- The switch port is capable of 4, 8, or 16 GB.
- The connected HBA maximum speed is 16 GB.
The highest common speed in this case is 16 GB, so the port should be configured for a speed of 16 GB.

```
FC_switch_A_1:admin> portcfgspeed 6 16
```

```
FC_switch_B_1:admin> portcfgspeed 6 16
```

2. Verify the settings:

```
portcfgshow
```

```
FC_switch_A_1:admin> portcfgshow
```

```
FC_switch_B_1:admin> portcfgshow
```

In the example output, port 6 has the following settings; speed is set to 16G:

Ports of Slot 0	0	1	2	3	4	5	6	7	8
Speed	16G								
AL_PA Offset 13
Trunk Port
Long Distance
VC Link Init
Locked L_Port	-	-	-	-	-	-	-	-	-
Locked G_Port
Disabled E_Port
Locked E_Port
ISL R_RDY Mode
RSCN Suppressed
Persistent Disable
LOS TOV enable
NPIV capability	ON								
NPIV PP Limit	126	126	126	126	126	126	126	126	126
QOS Port	AE	ON							
EX Port
Mirror Port
Rate Limit
Credit Recovery	ON								
Fport Buffers
Eport Credits
Port Auto Disable
CSCTL mode
D-Port mode
D-Port over DWDM
FEC	ON								
Fault Delay	0	0	0	0	0	0	0	0	0
Non-DFE

Configure compression on ISL ports on a Brocade G620 switch

If you are using Brocade G620 switches and enabling compression on the ISLs, you must configure it on each E-port on the switches.

About this task

This task must be performed on the ISL ports on both switches using the ISL.

Steps

1. Disable the port on which you want to configure compression:

```
portdisable port-id
```

2. Enable compression on the port:

```
portCfgCompress --enable port-id
```

3. Enable the port to activate the configuration with compression:

```
portenable port-id
```

4. Confirm that the setting has been changed:

```
portcfgshow port-id
```

The following example enables compression on port 0.

```
FC_switch_A_1:admin> portdisable 0
FC_switch_A_1:admin> portcfgcompress --enable 0
FC_switch_A_1:admin> portenable 0
FC_switch_A_1:admin> portcfgshow 0
Area Number: 0
Octet Speed Combo: 3(16G,10G)
(output truncated)
D-Port mode: OFF
D-Port over DWDM ..
Compression: ON
Encryption: ON
```

You can use the `islshow` command to check that the E_port has come online with encryption or compression configured and active.

```
FC_switch_A_1:admin> islshow
1: 0-> 0 10:00:c4:f5:7c:8b:29:86    5 FC_switch_B_1
sp: 16.000G bw: 16.000G TRUNK QOS CR_RECOV ENCRYPTION COMPRESSION
```

You can use the `portEncCompShow` command to see which ports are active. In this example you can see that encryption and compression are configured and active on port 0.

```
FC_switch_A_1:admin> portenccompshow
User          Encryption          Compression          Config
Port  Configured  Active  Configured  Active  Speed
----  -
0     Yes        Yes     Yes         Yes    16G
```

Configure zoning on Brocade FC switches

You must assign the switch ports to separate zones to separate controller and storage traffic.

Zone the FC-VI ports

For each DR group in the MetroCluster, you must configure two zones for the FC-VI connections that allow controller-to-controller traffic. These zones contain the FC switch ports connecting to the controller module FC-VI ports. These zones are Quality of Service (QoS) zones.

A QoS zone name starts with the prefix QOSHid_, followed by a user-defined string to differentiate it from a regular zone. These QoS zones are the same regardless of the model of FibreBridge bridge that is being used.

Each zone contains all the FC-VI ports, one for each FC-VI cable from each controller. These zones are configured for high priority.

The following tables show the FC-VI zones for two DR groups.

DR group 1 : QOSH1 FC-VI zone for FC-VI port a / c

FC switch	Site	Switch domain	6505 / 6510 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	0	0	0	controller_A_1 port FC-VI a
FC_switch_A_1	A	5	1	1	1	controller_A_1 port FC-VI c
FC_switch_A_1	A	5	4	4	4	controller_A_2 port FC-VI a
FC_switch_A_1	A	5	5	5	5	controller_A_2 port FC-VI c
FC_switch_B_1	B	7	0	0	0	controller_B_1 port FC-VI a
FC_switch_B_1	B	7	1	1	1	controller_B_1 port FC-VI c
FC_switch_B_1	B	7	4	4	4	controller_B_2 port FC-VI a
FC_switch_B_1	B	7	5	5	5	controller_B_2 port FC-VI c

Zone in Fabric_1	Member ports
QOSH1_MC1_FAB_1_FCVI	5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5

DR group 1 : QOSH1 FC-VI zone for FC-VI port b / d

FC switch	Site	Switch domain	6505 / 6510 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	0	0	0	controller_A_1 port FC-VI b
			1	1	1	controller_A_1 port FC-VI d

FC switch	Site	Switch domain	6505 / 6510 port	6520 port	G620 port	Connects to...
			4	4	4	controller_A_2 port FC-VI b
			5	5	5	controller_A_2 port FC-VI d
FC_switch_B_2	B	8	0	0	0	controller_B_1 port FC-VI b
			1	1	1	controller_B_1 port FC-VI d
			4	4	4	controller_B_2 port FC-VI b
			5	5	5	controller_B_2 port FC-VI d

Zone in Fabric_1	Member ports
QOSH1_MC1_FAB_2_FCVI	6,0;6,1;6,4;6,5;8,0;8,1;8,4;8,5

DR group 2 : QOSH2 FC-VI zone for FC-VI port a / c

FC switch	Site	Switch domain	Switch port			Connects to...
			6510	6520	G620	
FC_switch_A_1	A	5	24	48	18	controller_A_3 port FC-VI a
			25	49	19	controller_A_3 port FC-VI c
			28	52	22	controller_A_4 port FC-VI a
			29	53	23	controller_A_4 port FC-VI c
FC_switch_B_1	B	7	24	48	18	controller_B_3 port FC-VI a
			25	49	19	controller_B_3 port FC-VI c
			28	52	22	controller_B_4 port FC-VI a
			29	53	23	controller_B_4 port FC-VI c

Zone in Fabric_1	Member ports
QOSH2_MC2_FAB_1_FCVI (6510)	5,24;5,25;5,28;5,29;7,24;7,25;7,28;7,29

QOSH2_MC2_FAB_1_FCVI (6520)	5,48;5,49;5,52;5,53;7,48;7,49;7,52;7,53
-----------------------------	---

DR group 2 : QOSH2 FC-VI zone for FC-VI port b / d

FC switch	Site	Switch domain	6510 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	24	48	18	controller_A_3 port FC-VI b
FC_switch_A_2	A	6	25	49	19	controller_A_3 port FC-VI d
FC_switch_A_2	A	6	28	52	22	controller_A_4 port FC-VI b
FC_switch_A_2	A	6	29	53	23	controller_A_4 port FC-VI d
FC_switch_B_2	B	8	24	48	18	controller_B_3 port FC-VI b
FC_switch_B_2	B	8	25	49	19	controller_B_3 port FC-VI d
FC_switch_B_2	B	8	28	52	22	controller_B_4 port FC-VI b
FC_switch_B_2	B	8	29	53	23	controller_B_4 port FC-VI d

Zone in Fabric_2	Member ports
QOSH2_MC2_FAB_2_FCVI (6510)	6,24;6,25;6,28;6,29;8,24;8,25;8,28;8,29
QOSH2_MC2_FAB_2_FCVI (6520)	6,48;6,49;6,52;6,53;8,48;8,49;8,52;8,53

The following table provides a summary of the FC-VI zones:

Fabric	Zone name	Member ports
FC_switch_A_1 and FC_switch_B_1	QOSH1_MC1_FAB_1_FCVI	5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5
	QOSH2_MC1_FAB_1_FCVI (6510)	5,24;5,25;5,28;5,29;7,24;7,25;7,28;7,29
	QOSH2_MC1_FAB_1_FCVI (6520)	5,48;5,49;5,52;5,53;7,48;7,49;7,52;7,53

FC_switch_A_2 and FC_switch_B_2	QOSH1_MC1_FAB_2_FCVI	6,0;6,1;6,4;6,5;8,0;8,1;8,4;8,5
	QOSH2_MC1_FAB_2_FCVI (6510)	6,24;6,25;6,28;6,29;8,24;8,25;8,28; 8,29
	QOSH2_MC1_FAB_2_FCVI (6520)	6,48;6,49;6,52;6,53;8,48;8,49;8,52; 8,53

Zone FibreBridge 7500N or 7600N bridges using one FC port

If you are using FibreBridge 7500N or 7600N bridges using only one of the two FC ports, you need to create storage zones for the bridge ports. You should understand the zones and associated ports before you configure the zones.

The examples show zoning for DR group 1 only. If your configuration includes a second DR group, configure the zoning for the second DR group in the same manner, using the corresponding ports of the controllers and bridges.

Required zones

You must configure one zone for each of the FC-to-SAS bridge FC ports that allows traffic between initiators on each controller module and that FC-to-SAS bridge.

Each storage zone contains nine ports:

- Eight HBA initiator ports (two connections for each controller)
- One port connecting to an FC-to-SAS bridge FC port

The storage zones use standard zoning.

The examples show two pairs of bridges connecting two stack groups at each site. Because each bridge uses one FC port, there are a total of four storage zones per fabric (eight in total).

Bridge naming

The bridges use the following example naming: bridge_site_stack grouplocation in pair

This portion of the name...	Identifies the...	Possible values...
site	Site on which the bridge pair physically resides.	A or B
stack group	Number of the stack group to which the bridge pair connects. FibreBridge 7600N or 7500N bridges support up to four stacks in the stack group. The stack group can contain no more than 10 storage shelves.	1, 2, etc.

location in pair	Bridge within the bridge pair. A pair of bridges connect to a specific stack group.	a or b
------------------	---	--------

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

DR Group 1 - Stack 1 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_A_1	A	5	8	bridge_A_1a FC1
FC_switch_B_1	B	7	2	controller_B_1 port 0a
FC_switch_B_1	B	7	3	controller_B_1 port 0c
FC_switch_B_1	B	7	6	controller_B_2 port 0a
FC_switch_B_1	B	7	7	controller_B_2 port 0c

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,8

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d
FC_switch_A_1	A	6	8	bridge_A_1b FC1
FC_switch_B_1	B	8	2	controller_B_1 port 0b
FC_switch_B_1	B	8	3	controller_B_1 port 0d
FC_switch_B_1	B	8	6	controller_B_2 port 0b
FC_switch_B_1	B	8	7	controller_B_2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,8

DR Group 1 - Stack 2 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_A_1	A	5	9	bridge_A_2a FC1

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_B_1	B	7	2	controller_B_1 port 0a
FC_switch_B_1	B	7	3	controller_B_1 port 0c
FC_switch_B_1	B	7	6	controller_B_2 port 0a
FC_switch_B_1	B	7	7	controller_B_2 port 0c

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d
FC_switch_A_1	A	6	9	bridge_A_2b FC1
FC_switch_B_1	B	8	2	controller_B_1 port 0b
FC_switch_B_1	B	8	3	controller_B_1 port 0d
FC_switch_B_1	B	8	6	controller_B_2 port 0b
FC_switch_B_1	B	8	7	controller_B_2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,9

DR Group 1 - Stack 1 at Site_B

MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch	Connects to...
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_B_1	B	7	2	controller_B_1 port 0a
FC_switch_B_1	B	7	3	controller_B_1 port 0c
FC_switch_B_1	B	7	6	controller_B_2 port 0a
FC_switch_B_1	B	7	7	controller_B_2 port 0c
FC_switch_B_1	B	7	8	bridge_B_1a FC1

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,8

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_1_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch	Connects to...
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch	Connects to...
FC_switch_B_1	B	8	2	controller_B_1 port 0b
FC_switch_B_1	B	8	3	controller_B_1 port 0d
FC_switch_B_1	B	8	6	controller_B_2 port 0b
FC_switch_B_1	B	8	7	controller_B_2 port 0d
FC_switch_B_1	B	8	8	bridge_B_1b FC1

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_BOT_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;8,8

DR Group 1 - Stack 2 at Site_B

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_A_1	A	5	2	controller_A_1 port 0a
FC_switch_A_1	A	5	3	controller_A_1 port 0c
FC_switch_A_1	A	5	6	controller_A_2 port 0a
FC_switch_A_1	A	5	7	controller_A_2 port 0c
FC_switch_B_1	B	7	2	controller_B_1 port 0a
FC_switch_B_1	B	7	3	controller_B_1 port 0c
FC_switch_B_1	B	7	6	controller_B_2 port 0a
FC_switch_B_1	B	7	7	controller_B_2 port 0c
FC_switch_B_1	B	7	9	bridge_b_2a FC1

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_b_STK_GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_BOT_FC1:

FC switch	Site	Switch domain	Brocade 6505, 6510, 6520, G620, G610, or G710 switch port	Connects to...
FC_switch_A_1	A	6	2	controller_A_1 port 0b
FC_switch_A_1	A	6	3	controller_A_1 port 0d
FC_switch_A_1	A	6	6	controller_A_2 port 0b
FC_switch_A_1	A	6	7	controller_A_2 port 0d
FC_switch_B_1	B	8	2	controller_B_1 port 0b
FC_switch_B_1	B	8	3	controller_B_1 port 0d
FC_switch_B_1	B	8	6	controller_B_2 port 0b
FC_switch_B_1	B	8	7	controller_B_2 port 0d
FC_switch_B_1	B	8	9	bridge_B_1b FC1

Zone in Fabric_2	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;8,9

Summary of storage zones

Fabric	Zone name	Member ports
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,8
	MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;5,9
	MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,8
	MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1	5,2;5,3;5,6;5,7;7,2;7,3;7,6;7,7;7,9

FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,8
	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;6,9
	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;8,8
	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_BOT_FC1	6,2;6,3;6,6;6,7;8,2;8,3;8,6;8,7;8,9

Zone FibreBridge 7500N bridges using both FC ports

If you are using FibreBridge 7500N bridges with both FC ports, you need to create storage zones for the bridge ports. You should understand the zones and associated ports before you configure the zones.

Required zones

You must configure one zone for each of the FC-to-SAS bridge FC ports that allows traffic between initiators on each controller module and that FC-to-SAS bridge.

Each storage zone contains five ports:

- Four HBA initiator ports (one connection for each controller)
- One port connecting to an FC-to-SAS bridge FC port

The storage zones use standard zoning.

The examples show two pairs of bridges connecting two stack groups at each site. Because each bridge uses one FC port, there are a total of eight storage zones per fabric (sixteen in total).

Bridge naming

The bridges use the following example naming: bridge_site_stack grouplocation in pair

This portion of the name...	Identifies the...	Possible values...
site	Site on which the bridge pair physically resides.	A or B
stack group	Number of the stack group to which the bridge pair connects. FibreBridge 7600N or 7500N bridges support up to four stacks in the stack group. The stack group can contain no more than 10 storage shelves.	1, 2, etc.

location in pair	Bridge within the bridge pair. A pair of bridges connect to a specific stack group.	a or b
------------------	---	--------

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

DR Group 1 - Stack 1 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 / G620 port	6520 port	Connects to...
FC_switch_A_1	A	5	2	2	controller_A_1 port 0a
FC_switch_A_1	A	5	6	6	controller_A_2 port 0a
FC_switch_A_1	A	5	8	8	bridge_A_1a FC1
FC_switch_B_1	B	7	2	2	controller_B_1 port 0a
FC_switch_B_1	B	7	6	6	controller_B_2 port 0a

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;5,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	3	3	3	controller_A_1 port 0c

FC_switch_A_1	A	5	7	7	7	controller_A_2 port 0c
FC_switch_A_1	A	5	9	9	9	bridge_A_1b FC1
FC_switch_B_1	B	7	3	3	3	controller_B_1 port 0c
FC_switch_B_1	B	7	7	7	7	controller_B_2 port 0c

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;5,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_1_BOT_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710	6520	G620	Connects to...
FC_switch_A_2	A	6	2	2	2	controller_A_1 port 0b
FC_switch_A_2	A	6	6	6	6	controller_A_2 port 0b
FC_switch_A_2	A	6	8	8	8	bridge_A_1a FC2
FC_switch_B_2	B	8	2	2	2	controller_B_1 port 0b
FC_switch_B_2	B	8	6	6	6	controller_B_2 port 0b

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;6,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_1_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710	6520	G620	Connects to...
FC_switch_A_2	A	6	3	3	3	controller_A_1 port 0d

FC_switch_A_2	A	6	7	7	7	controller_A_2 port 0d
FC_switch_A_2	A	6	9	9	9	bridge_A_1b FC2
FC_switch_B_2	B	8	3	3	3	controller_B_1 port 0d
FC_switch_B_2	B	8	7	7	7	controller_B_2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;6,9

DR Group 1 - Stack 2 at Site_A

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	2	2	2	controller_A_1 port 0a
FC_switch_A_1	A	5	6	6	6	controller_A_2 port 0a
FC_switch_A_1	A	5	10	10	10	bridge_A_2a FC1
FC_switch_B_1	B	7	2	2	2	controller_B_1 port 0a
FC_switch_B_1	B	7	6	6	6	controller_B_2 port 0a

Zone in Fabric_1 hh	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;5,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	3	3	3	controller_A_1 port 0c
FC_switch_A_1	A	5	7	7	7	controller_A_2 port 0c
FC_switch_A_1	A	5	11	11	11	bridge_A_2b FC1
FC_switch_B_1	B	7	3	3	3	controller_B_1 port 0c
FC_switch_B_1	B	7	7	7	7	controller_B_2 port 0c

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;5,11

DrGroup 1 : MC1_INIT_GRP_1_SITE_A_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	2	0	0	controller_A_1 port 0b
FC_switch_A_2	A	6	6	4	4	controller_A_2 port 0b
FC_switch_A_2	A	6	10	10	10	bridge_A_2a FC2
FC_switch_B_2	B	8	2	2	2	controller_B_1 port 0b
FC_switch_B_2	B	8	6	6	6	controller_B_2 port 0b

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_A_STK_GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;6,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_A_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	3	3	3	controller_A_1 port 0d
FC_switch_A_2	A	6	7	7	7	controller_A_2 port 0d
FC_switch_A_2	A	6	11	11	11	bridge_A_2b FC2
FC_switch_B_2	B	8	3	3	3	controller_B_1 port 0d
FC_switch_B_2	B	8	7	7	7	controller_B_2 port 0d

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_A_STK_GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;6,11

DR Group 1 - Stack 1 at Site_B

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	2	2	2	controller_A_1 port 0a
FC_switch_A_1	A	5	6	6	6	controller_A_2 port 0a
FC_switch_B_1	B	7	2	2	8	controller_B_1 port 0a
FC_switch_B_1	B	7	6	6	2	controller_B_2 port 0a
FC_switch_B_1	B	7	8	8	6	bridge_B_1a FC1

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;7,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_1_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	3	3	3	controller_A_1 port 0c
FC_switch_A_1	A	5	7	7	7	controller_A_2 port 0c
FC_switch_B_1	B	7	3	3	9	controller_B_1 port 0c
FC_switch_B_1	B	7	7	7	3	controller_B_2 port 0c
FC_switch_B_1	B	7	9	9	7	bridge_B_1b FC1

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;7,9

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_1_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	2	2	2	controller_A_1 port 0b
FC_switch_A_2	A	6	6	6	6	controller_A_2 port 0b
FC_switch_B_2	B	8	2	2	2	controller_B_1 port 0b
FC_switch_B_2	B	8	6	6	6	controller_B_2 port 0b

FC_switch_B_2	B	8	8	8	8	bridge_B_1a FC2
---------------	---	---	---	---	---	--------------------

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;8,8

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_1_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	3	3	3	controller_A_1 port 0d
FC_switch_A_2	A	6	7	7	7	controller_A_2 port 0d
FC_switch_B_2	B	8	3	3	3	controller_B_1 port 0d
FC_switch_B_2	B	8	7	7	7	controller_B_2 port 0d
FC_switch_B_2	B	8	9	9	9	bridge_A_1b FC2

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;8,9

DR Group 1 - Stack 2 at Site_B

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	2	2	2	controller_A_1 port 0a
FC_switch_A_1	A	5	6	6	6	controller_A_2 port 0a

FC_switch_B_1	B	7	2	2	2	controller_B_1 port 0a
FC_switch_B_1	B	7	6	6	6	controller_B_2 port 0a
FC_switch_B_1	B	7	10	10	10	bridge_B_2a FC1

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;7,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_2_TOP_FC1:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_1	A	5	3	3	3	controller_A_1 port 0c
FC_switch_A_1	A	5	7	7	7	controller_A_2 port 0c
FC_switch_B_1	B	7	3	3	3	controller_B_1 port 0c
FC_switch_B_1	B	7	7	7	7	controller_B_2 port 0c
FC_switch_B_1	B	7	11	11	11	bridge_B_2b FC1

Zone in Fabric_2 hh	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;7,11

DrGroup 1 : MC1_INIT_GRP_1_SITE_B_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	2	2	2	controller_A_1 port 0b

FC_switch_A_2	A	6	6	6	6	controller_A_2 port 0b
FC_switch_B_2	B	8	2	2	2	controller_B_1 port 0b
FC_switch_B_2	B	8	6	6	6	controller_B_2 port 0b
FC_switch_B_2	B	8	10	10	10	bridge_B_2a FC2

Zone in Fabric_1	Member ports
MC1_INIT_GRP_1_SITE_B_STK_GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;8,10

DrGroup 1 : MC1_INIT_GRP_2_SITE_B_STK_GRP_2_BOT_FC2:

FC switch	Site	Switch domain	6505 / 6510 / G610 / G710 port	6520 port	G620 port	Connects to...
FC_switch_A_2	A	6	3	3	3	controller_A_1 port 0d
FC_switch_A_2	A	6	7	7	7	controller_A_2 port 0d
FC_switch_B_2	B	8	3	3	3	controller_B_1 port 0d
FC_switch_B_2	B	8	7	7	7	controller_B_2 port 0d
FC_switch_B_2	B	8	11	11	11	bridge_B_2b FC2

Zone in Fabric_2	Member ports
MC1_INIT_GRP_2_SITE_B_STK_GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;8,11

Summary of storage zones

Fabric	Zone name	Member ports
--------	-----------	--------------

FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;5,8
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_A_STK_ GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;5,9
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;5,10
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_A_STK_ GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;5,11
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_TOP_FC1	5,2;5,6;7,2;7,6;7,8
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_B_STK_ GRP_1_BOT_FC1	5,3;5,7;7,3;7,7;7,9
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_TOP_FC1	5,2;5,6;7,2;7,6;7,10
FC_switch_A_1 and FC_switch_B_1	MC1_INIT_GRP_2_SITE_B_STK_ GRP_2_BOT_FC1	5,3;5,7;7,3;7,7;7,11
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_A_STK_ GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;6,8
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_A_STK_ GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;6,9
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_A_STK_ GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;6,10
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_A_STK_ GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;6,11
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_B_STK_ GRP_1_TOP_FC2	6,2;6,6;8,2;8,6;8,8
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_B_STK_ GRP_1_BOT_FC2	6,3;6,7;8,3;8,7;8,9
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_1_SITE_B_STK_ GRP_2_TOP_FC2	6,2;6,6;8,2;8,6;8,10
FC_switch_A_2 and FC_switch_B_2	MC1_INIT_GRP_2_SITE_B_STK_ GRP_2_BOT_FC2	6,3;6,7;8,3;8,7;8,11

Zone the Brocade FC switches

You must assign the switch ports to separate zones to separate controller and storage traffic, with zones for the FC-VI ports and zones for the storage ports.

About this task

The following steps use the standard zoning for the MetroCluster configuration.

Zoning for FC-VI ports

Zoning for FibreBridge 7500N or 7600N bridges using one FC port

Zoning for FibreBridge 7500N bridges using both FC ports

Steps

1. Create the FC-VI zones on each switch:

```
zonecreate "QOSH1_FCVI_1", member;member ...
```

In this example a QOS FCVI zone is created containing ports 5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5:

```
Switch_A_1:admin> zonecreate "QOSH1_FCVI_1",  
"5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5"
```

2. Configure the storage zones on each switch.

You can configure zoning for the fabric from one switch in the fabric. In the example that follows, zoning is configured on Switch_A_1.

- a. Create the storage zone for each switch domain in the switch fabric:

```
zonecreate name, member;member ...
```

In this example a storage zone for a FibreBridge 7500N using both FC ports is being created. The zones contains ports 5,2;5,6;7,2;7,6;5,16:

```
Switch_A_1:admin> zonecreate  
"MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1", "5,2;5,6;7,2;7,6;5,16"
```

- b. Create the configuration in the first switch fabric:

```
cfgcreate config_name, zone;zone...
```

In this example a configuration with the name CFG_1 and the two zones QOSH1_MC1_FAB_1_FCVI and MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1 is created

```
Switch_A_1:admin> cfgcreate "CFG_1", "QOSH1_MC1_FAB_1_FCVI;  
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1"
```

c. Add zones to the configuration, if desired:

```
cfgadd config_namezone;zone...
```

d. Enable the configuration:

```
cfgenable config_name
```

```
Switch_A_1:admin> cfgenable "CFG_1"
```

e. Save the configuration:

```
cfgsave
```

```
Switch_A_1:admin> cfgsave
```

f. Validate the zoning configuration:

```
zone --validate
```

```

Switch_A_1:admin> zone --validate
Defined configuration:
cfg: CFG_1 QOSH1_MC1_FAB_1_FCVI ;
MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1
zone: QOSH1_MC1_FAB_1_FCVI
5,0;5,1;5,4;5,5;7,0;7,1;7,4;7,5
zone: MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1
5,2;5,6;7,2;7,6;5,16
Effective configuration:
cfg: CFG_1
zone: QOSH1_MC1_FAB_1_FCVI
5,0
5,1
5,4
5,5
7,0
7,1
7,4
7,5
zone: MC1_INIT_GRP_1_SITE_A_STK_GRP_1_TOP_FC1
5,2
5,6
7,2
7,6
5,16
-----
~ - Invalid configuration
* - Member does not exist
# - Invalid usage of broadcast zone

```

Set ISL encryption on Brocade 6510 or G620 switches

On Brocade 6510 or G620 switches, you can optionally use the Brocade encryption feature on the ISL connections. If you want to use the encryption feature, you must perform additional configuration steps on each switch in the MetroCluster configuration.

Before you begin

- You must have Brocade 6510 or G620 switches.



Support for ISL encryption on Brocade G620 switches is only supported on ONTAP 9.4 and later.

- You must have selected two switches from the same fabric.
- You must have reviewed the Brocade documentation for your switch and Fabric Operating System version to confirm the bandwidth and port limits.

About this task

The steps must be performed on both the switches in the same fabric.

Disable virtual fabric

In order to set the ISL encryption, you must disable the virtual fabric on all the four switches being used in a MetroCluster configuration.

Steps

1. Disable the virtual fabric by entering the following command at the switch console:

```
fosconfig --disable vf
```

2. Reboot the switch.

Set the payload

After disabling the virtual fabric, you must set the payload or the data field size on both switches in the fabric.

About this task

The data field size must not exceed 2048.

Steps

1. Disable the switch:

```
switchdisable
```

2. Configure and set the payload:

```
configure
```

3. Set the following switch parameters:
 - a. Set the Fabric parameter as follows: `y`
 - b. Set the other parameters, such as Domain, WWN-based persistent PID, and so on.
 - c. Set the data field size: `2048`

Set the authentication policy

You must set the authentication policy and associated parameters.

About this task

The commands must be executed at the switch console.

Steps

1. Set the authentication secret:
 - a. Begin the setup process:

```
secAuthSecret --set
```

This command initiates a series of prompts that you respond to in the following steps:

- b. Provide the worldwide name (WWN) of the other switch in the fabric for the "Enter peer WWN, Domain, or switch name" parameter.
- c. Provide the peer secret for the "Enter peer secret" parameter.
- d. Provide the local secret for the "Enter local secret" parameter.
- e. Enter Y for the "Are you done" parameter.

The following is an example of setting the authentication secret:

```
brcd> secAuthSecret --set
```

This command is used to set up secret keys for the DH-CHAP authentication.

The minimum length of a secret key is 8 characters and maximum 40 characters. Setting up secret keys does not initiate DH-CHAP authentication. If switch is configured to do DH-CHAP, it is performed whenever a port or a switch is enabled.

Warning: Please use a secure channel for setting secrets. Using an insecure channel is not safe and may compromise secrets.

Following inputs should be specified for each entry.

1. WWN for which secret is being set up.
2. Peer secret: The secret of the peer that authenticates to peer.
3. Local secret: The local secret that authenticates peer.

Press enter to start setting up secrets > <cr>

Enter peer WWN, Domain, or switch name (Leave blank when done):

10:00:00:05:33:76:2e:99

Enter peer secret: <hidden>

Re-enter peer secret: <hidden>

Enter local secret: <hidden>

Re-enter local secret: <hidden>

Enter peer WWN, Domain, or switch name (Leave blank when done):

Are you done? (yes, y, no, n): [no] yes

Saving data to key store... Done.

2. Set the authentication group to 4:

```
authUtil --set -g 4
```

3. Set the authentication type to "dhchap":

```
authUtil --set -a dhchap
```

The system displays the following output:

```
Authentication is set to dhchap.
```

4. Set the authentication policy on the switch to on:

```
authUtil --policy -sw on
```

The system displays the following output:

```
Warning: Activating the authentication policy requires either DH-CHAP
secrets or PKI certificates depending on the protocol selected.
Otherwise, ISLs will be segmented during next E-port bring-up.
ARE YOU SURE (yes, y, no, n): [no] yes
Auth Policy is set to ON
```

Enable ISL encryption on Brocade switches

After setting the authentication policy and the authentication secret, you must enable ISL encryption on the ports for it to take effect.

About this task

- These steps should be performed on one switch fabric at a time.
- The commands must be run at the switch console.

Steps

1. Enable encryption on all of the ISL ports:

```
portCfgEncrypt --enable port_number
```

In the following example, the encryption is enabled on ports 8 and 12:

```
portCfgEncrypt --enable 8
```

```
portCfgEncrypt --enable 12
```

2. Enable the switch:

```
switchenable
```

3. Verify that the ISL is up and working:

```
islshow
```

4. Verify that encryption is enabled:

```
portenccompshow
```

The following example shows that encryption is enabled on ports 8 and 12:

User Port	Encryption configured	Active
8	yes	yes
9	No	No
10	No	No
11	No	No
12	yes	yes

What to do next

Perform all of the steps on the switches in the other fabric in a MetroCluster configuration.

Configuring the Cisco FC switches manually

Each Cisco switch in the MetroCluster configuration must be configured appropriately for the ISL and storage connections.

Before you begin

The following requirements apply to the Cisco FC switches:

- You must use four supported Cisco switches of the same model with the same NX-OS version and licensing.
- The MetroCluster configuration requires four switches.

The four switches must be connected into two fabrics of two switches each, with each fabric spanning both sites.

- The switch must support connectivity to the ATTO FibreBridge model.
- You cannot use encryption or compression in the Cisco FC storage fabric. It is not supported in the MetroCluster configuration.

In the [NetApp Interoperability Matrix Tool \(IMT\)](#), you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

About this task

The following requirement applies to the Inter-Switch Link (ISL) connections:

- All ISLs must have the same length and same speed in one fabric.

Different lengths of ISLs can be used in the different fabrics. The same speed must be used in all fabrics.

The following requirement applies to the storage connections:

- Each storage controller must have four initiator ports available to connect to the switch fabrics.

Two initiator ports must be connected from each storage controller to each fabric.



You can configure FAS8020, AFF8020, FAS8200, and AFF A300 systems with two initiators ports per controller (a single initiator port to each fabric) if all of the following criteria are met:

- There are fewer than four FC initiator ports available to connect the disk storage and no additional ports can be configured as FC initiators.
- All slots are in use and no FC initiator card can be added.

Related information

[NetApp Interoperability Matrix Tool](#)

Cisco switch license requirements

Certain feature-based licenses might be required for the Cisco switches in a fabric-attached MetroCluster configuration. These licenses enable you to use features such as QoS or long-distance mode credits on the switches. You must install the required feature-based licenses on all four switches in a MetroCluster configuration.

The following feature-based licenses might be required in a MetroCluster configuration:

- ENTERPRISE_PKG

This license enables you to use the QoS feature on Cisco switches.

- PORT_ACTIVATION_PKG

You can use this license for Cisco 9148 switches. This license enables you to activate or deactivate ports on the switches as long as only 16 ports are active at any given time. By default, 16 ports are enabled in Cisco MDS 9148 switches.

- FM_SERVER_PKG

This license enables you to manage fabrics simultaneously and to manage switches through a web browser.

The FM_SERVER_PKG license also enables performance management features such as performance thresholds and threshold monitoring. For more information about this license, see the Cisco Fabric Manager Server Package.

You can verify that the licenses are installed by using the show license usage command. If you do not have these licenses, contact your sales representative before proceeding with the installation.



The Cisco MDS 9250i switches have two fixed 1/10 GbE IP storage services ports. No additional licenses are required for these ports. The Cisco SAN Extension over IP application package is a standard license on these switches that enables features such as FCIP and compression.

Setting the Cisco FC switch to factory defaults

To ensure a successful configuration, you must set the switch to its factory defaults. This ensures that the switch is starting from a clean configuration.

About this task

This task must be performed on all switches in the MetroCluster configuration.

Steps

1. Make a console connection and log in to both switches in the same fabric.
2. Set the switch back to its default settings:

```
write erase
```

You can respond “y” when prompted to confirm the command. This erases all licenses and configuration information on the switch.

3. Reboot the switch:

```
reload
```

You can respond “y” when prompted to confirm the command.

4. Repeat the `write erase` and `reload` commands on the other switch.

After issuing the `reload` command, the switch reboots and then prompts with setup questions. At that point, proceed to the next section.

Example

The following example shows the process on a fabric consisting of FC_switch_A_1 and FC_switch_B_1.

```
FC_Switch_A_1# write erase
Warning: This command will erase the startup-configuration.
Do you wish to proceed anyway? (y/n) [n] y
FC_Switch_A_1# reload
This command will reboot the system. (y/n)? [n] y

FC_Switch_B_1# write erase
Warning: This command will erase the startup-configuration.
Do you wish to proceed anyway? (y/n) [n] y
FC_Switch_B_1# reload
This command will reboot the system. (y/n)? [n] y
```

Configure the Cisco FC switch basic settings and community string

You must specify the basic settings with the `setup` command or after issuing the `reload` command.

Steps

1. If the switch does not display the setup questions, configure the basic switch settings:

```
setup
```

2. Accept the default responses to the setup questions until you are prompted for the SNMP community string.

3. Set the community string to “public” (all lowercase) to allow access from the ONTAP Health Monitors.

You can set the community string to a value other than “public”, but you must configure the ONTAP Health Monitors using the community string you specify.

The following example shows the commands on FC_switch_A_1:

```
FC_switch_A_1# setup
  Configure read-only SNMP community string (yes/no) [n]: y
  SNMP community string : public
  Note: Please set the SNMP community string to "Public" or another
value of your choosing.
  Configure default switchport interface state (shut/noshut) [shut]:
noshut
  Configure default switchport port mode F (yes/no) [n]: n
  Configure default zone policy (permit/deny) [deny]: deny
  Enable full zoneset distribution? (yes/no) [n]: yes
```

The following example shows the commands on FC_switch_B_1:

```
FC_switch_B_1# setup
  Configure read-only SNMP community string (yes/no) [n]: y
  SNMP community string : public
  Note: Please set the SNMP community string to "Public" or another
value of your choosing.
  Configure default switchport interface state (shut/noshut) [shut]:
noshut
  Configure default switchport port mode F (yes/no) [n]: n
  Configure default zone policy (permit/deny) [deny]: deny
  Enable full zoneset distribution? (yes/no) [n]: yes
```

Acquiring licenses for ports

You do not have to use Cisco switch licenses on a continuous range of ports; instead, you can acquire licenses for specific ports that are used and remove licenses from unused ports.

Before you begin

You should verify the number of licensed ports in the switch configuration and, if necessary, move licenses from one port to another as needed.

Steps

1. Display the license usage for a switch fabric:

```
show port-resources module 1
```

Determine which ports require licenses. If some of those ports are unlicensed, determine if you have extra licensed ports and consider removing the licenses from them.

2. Enter configuration mode:

```
config t
```

3. Remove the license from the selected port:

a. Select the port to be unlicensed:

```
interface interface-name
```

b. Remove the license from the port:

```
no port-license acquire
```

c. Exit the port configuration interface:

```
exit
```

4. Acquire the license for the selected port:

a. Select the port to be unlicensed:

```
interface interface-name
```

b. Make the port eligible to acquire a license:

```
port-license
```

c. Acquire the license on the port:

```
port-license acquire
```

d. Exit the port configuration interface:

```
exit
```

5. Repeat for any additional ports.

6. Exit configuration mode:

```
exit
```

Removing and acquiring a license on a port

This example shows a license being removed from port fc1/2, port fc1/1 being made eligible to acquire a license, and the license being acquired on port fc1/1:

```

Switch_A_1# conf t
Switch_A_1(config)# interface fc1/2
Switch_A_1(config)# shut
Switch_A_1(config-if)# no port-license acquire
Switch_A_1(config-if)# exit
Switch_A_1(config)# interface fc1/1
Switch_A_1(config-if)# port-license
Switch_A_1(config-if)# port-license acquire
Switch_A_1(config-if)# no shut
Switch_A_1(config-if)# end
Switch_A_1# copy running-config startup-config

Switch_B_1# conf t
Switch_B_1(config)# interface fc1/2
Switch_B_1(config)# shut
Switch_B_1(config-if)# no port-license acquire
Switch_B_1(config-if)# exit
Switch_B_1(config)# interface fc1/1
Switch_B_1(config-if)# port-license
Switch_B_1(config-if)# port-license acquire
Switch_B_1(config-if)# no shut
Switch_B_1(config-if)# end
Switch_B_1# copy running-config startup-config

```

The following example shows port license usage being verified:

```

Switch_A_1# show port-resources module 1
Switch_B_1# show port-resources module 1

```

Enabling ports in a Cisco MDS 9148 or 9148S switch

In Cisco MDS 9148 or 9148S switches, you must manually enable the ports required in a MetroCluster configuration.

About this task

- You can manually enable 16 ports in a Cisco MDS 9148 or 9148S switch.
- The Cisco switches enable you to apply the POD license on random ports, as opposed to applying them in sequence.
- Cisco switches require that you use one port from each port group, unless you need more than 12 ports.

Steps

1. View the port groups available in a Cisco switch:

```
show port-resources module blade_number
```

2. License and acquire the required port in a port group:

```
config t  
  
interface port_number  
  
shut  
  
port-license acquire  
  
no shut
```

For example, the following command sequence licenses and acquires Port fc 1/45:

```
switch# config t  
switch(config)#  
switch(config)# interface fc 1/45  
switch(config-if)#  
switch(config-if)# shut  
switch(config-if)# port-license acquire  
switch(config-if)# no shut  
switch(config-if)# end
```

3. Save the configuration:

```
copy running-config startup-config
```

Configuring the F-ports on a Cisco FC switch

You must configure the F-ports on the FC switch.

About this task

In a MetroCluster configuration, the F-ports are the ports that connect the switch to the HBA initiators, FC-VI interconnects and FC-to-SAS bridges.

Each port must be configured individually.

Refer to the following sections to identify the F-ports (switch-to-node) for your configuration:

- [Port assignments for FC switches](#)

This task must be performed on each switch in the MetroCluster configuration.

Steps

1. Enter configuration mode:

```
config t
```

2. Enter interface configuration mode for the port:

```
interface port-ID
```

3. Shut down the port:

```
shutdown
```

4. Set the ports to F mode:

```
switchport mode F
```

5. Set the ports to fixed speed:

```
switchport speed speed-value
```

speed-value is either 8000 or 16000

6. Set the rate mode of the switch port to dedicated:

```
switchport rate-mode dedicated
```

7. Restart the port:

```
no shutdown
```

8. Exit configuration mode:

```
end
```

Example

The following example shows the commands on the two switches:

```
Switch_A_1# config t
FC_switch_A_1(config)# interface fc 1/1
FC_switch_A_1(config-if)# shutdown
FC_switch_A_1(config-if)# switchport mode F
FC_switch_A_1(config-if)# switchport speed 8000
FC_switch_A_1(config-if)# switchport rate-mode dedicated
FC_switch_A_1(config-if)# no shutdown
FC_switch_A_1(config-if)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# config t
FC_switch_B_1(config)# interface fc 1/1
FC_switch_B_1(config-if)# switchport mode F
FC_switch_B_1(config-if)# switchport speed 8000
FC_switch_B_1(config-if)# switchport rate-mode dedicated
FC_switch_B_1(config-if)# no shutdown
FC_switch_B_1(config-if)# end
FC_switch_B_1# copy running-config startup-config
```

Assigning buffer-to-buffer credits to F-Ports in the same port group as the ISL

You must assign the buffer-to-buffer credits to the F-ports if they are in the same port group as the ISL. If the ports do not have the required buffer-to-buffer credits, the ISL could be inoperative.

About this task

This task is not required if the F-ports are not in the same port group as the ISL port.

If the F-Ports are in a port group that contains the ISL, this task must be performed on each FC switch in the MetroCluster configuration.

Steps

1. Enter configuration mode:

```
config t
```

2. Set the interface configuration mode for the port:

```
interface port-ID
```

3. Disable the port:

```
shut
```

4. If the port is not already in F mode, set the port to F mode:

```
switchport mode F
```

5. Set the buffer-to-buffer credit of the non-E ports to 1:

```
switchport fcrxbbcredit 1
```

6. Re-enable the port:

```
no shut
```

7. Exit configuration mode:

```
exit
```

8. Copy the updated configuration to the startup configuration:

```
copy running-config startup-config
```

9. Verify the buffer-to-buffer credit assigned to a port:

```
show port-resources module 1
```

10. Exit configuration mode:

```
exit
```

11. Repeat these steps on the other switch in the fabric.

12. Verify the settings:

```
show port-resource module 1
```

Example

In this example, port fc1/40 is the ISL. Ports fc1/37, fc1/38 and fc1/39 are in the same port group and must be configured.

The following commands show the port range being configured for fc1/37 through fc1/39:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# interface fc1/37-39
FC_switch_A_1(config-if)# shut
FC_switch_A_1(config-if)# switchport mode F
FC_switch_A_1(config-if)# switchport fcrxbbcredit 1
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config-if)# exit
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# conf t
FC_switch_B_1(config)# interface fc1/37-39
FC_switch_B_1(config-if)# shut
FC_switch_B_1(config-if)# switchport mode F
FC_switch_B_1(config-if)# switchport fcrxbbcredit 1
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config-if)# exit
FC_switch_B_1# copy running-config startup-config
```

The following commands and system output show that the settings are properly applied:

```

FC_switch_A_1# show port-resource module 1
...
Port-Group 11
  Available dedicated buffers are 93

-----
Interfaces in the Port-Group          B2B Credit  Bandwidth  Rate Mode
                                   Buffers      (Gbps)
-----
fc1/37                               32          8.0    dedicated
fc1/38                               1           8.0    dedicated
fc1/39                               1           8.0    dedicated
...

FC_switch_B_1# port-resource module
...
Port-Group 11
  Available dedicated buffers are 93

-----
Interfaces in the Port-Group          B2B Credit  Bandwidth  Rate Mode
                                   Buffers      (Gbps)
-----
fc1/37                               32          8.0    dedicated
fc1/38                               1           8.0    dedicated
fc1/39                               1           8.0    dedicated
...

```

Creating and configuring VSANs on Cisco FC switches

You must create a VSAN for the FC-VI ports and a VSAN for the storage ports on each FC switch in the MetroCluster configuration.

About this task

The VSANs should have a unique number and name. You must do additional configuration if you are using two ISLs with in-order delivery of frames.

The examples of this task use the following naming conventions:

Switch fabric	VSAN name	ID number
1	FCVI_1_10	10
	STOR_1_20	20

2	FCVI_2_30	30
	STOR_2_20	40

This task must be performed on each FC switch fabric.

Steps

1. Configure the FC-VI VSAN:

- a. Enter configuration mode if you have not done so already:

```
config t
```

- b. Edit the VSAN database:

```
vsan database
```

- c. Set the VSAN ID:

```
vsan vsan-ID
```

- d. Set the VSAN name:

```
vsan vsan-ID name vsan_name
```

2. Add ports to the FC-VI VSAN:

- a. Add the interfaces for each port in the VSAN:

```
vsan vsan-ID interface interface_name
```

For the FC-VI VSAN, the ports connecting the local FC-VI ports will be added.

- b. Exit configuration mode:

```
end
```

- c. Copy the running-config to the startup-config:

```
copy running-config startup-config
```

In the following example, the ports are fc1/1 and fc1/13:

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config)# vsan 10 interface fc1/1
FC_switch_A_1(config)# vsan 10 interface fc1/13
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config
FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config)# vsan 10 interface fc1/1
FC_switch_B_1(config)# vsan 10 interface fc1/13
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config

```

3. Verify port membership of the VSAN:

```
show vsan member
```

```

FC_switch_A_1# show vsan member
FC_switch_B_1# show vsan member

```

4. Configure the VSAN to guarantee in-order delivery of frames or out-of-order delivery of frames:



The standard IOD settings are recommended. You should configure OOD only if necessary.

Considerations for using TDM/WDM equipment with fabric-attached MetroCluster configurations

- The following steps must be performed to configure in-order delivery of frames:

- a. Enter configuration mode:

```
conf t
```

- b. Enable the in-order guarantee of exchanges for the VSAN:

```
in-order-guarantee vsan vsan-ID
```



For FC-VI VSANs (FCVI_1_10 and FCVI_2_30), you must enable in-order guarantee of frames and exchanges only on VSAN 10.

- c. Enable load balancing for the VSAN:

```
vsan vsan-ID loadbalancing src-dst-id
```

- d. Exit configuration mode:

```
end
```

- e. Copy the running-config to the startup-config:

```
copy running-config startup-config
```

The commands to configure in-order delivery of frames on FC_switch_A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# in-order-guarantee vsan 10
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to configure in-order delivery of frames on FC_switch_B_1:

```
FC_switch_B_1# config t
FC_switch_B_1(config)# in-order-guarantee vsan 10
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config
```

- The following steps must be performed to configure out-of-order delivery of frames:

- a. Enter configuration mode:

```
conf t
```

- b. Disable the in-order guarantee of exchanges for the VSAN:

```
no in-order-guarantee vsan vsan-ID
```

- c. Enable load balancing for the VSAN:

```
vsan vsan-ID loadbalancing src-dst-id
```

- d. Exit configuration mode:

```
end
```

- e. Copy the running-config to the startup-config:

```
copy running-config startup-config
```

The commands to configure out-of-order delivery of frames on FC_switch_A_1:

```

FC_switch_A_1# config t
FC_switch_A_1(config)# no in-order-guarantee vsan 10
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config

```

The commands to configure out-of-order delivery of frames on FC_switch_B_1:

```

FC_switch_B_1# config t
FC_switch_B_1(config)# no in-order-guarantee vsan 10
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config

```



When configuring ONTAP on the controller modules, OOD must be explicitly configured on each controller module in the MetroCluster configuration.

Configuring in-order delivery or out-of-order delivery of frames on ONTAP software

5. Set QoS policies for the FC-VI VSAN:

- a. Enter configuration mode:

```
conf t
```

- b. Enable the QoS and create a class map by entering the following commands in sequence:

```
qos enable
```

```
qos class-map class_name match-any
```

- c. Add the class map created in a previous step to the policy map:

```
class class_name
```

- d. Set the priority:

```
priority high
```

- e. Add the VSAN to the policy map created previously in this procedure:

```
qos service policy policy_name vsan vsan-id
```

- f. Copy the updated configuration to the startup configuration:

```
copy running-config startup-config
```

The commands to set the QoS policies on FC_switch_A_1:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# qos enable
FC_switch_A_1(config)# qos class-map FCVI_1_10_Class match-any
FC_switch_A_1(config)# qos policy-map FCVI_1_10_Policy
FC_switch_A_1(config-pmap)# class FCVI_1_10_Class
FC_switch_A_1(config-pmap-c)# priority high
FC_switch_A_1(config-pmap-c)# exit
FC_switch_A_1(config)# exit
FC_switch_A_1(config)# qos service policy FCVI_1_10_Policy vsan 10
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to set the QoS policies on FC_switch_B_1:

```
FC_switch_B_1# conf t
FC_switch_B_1(config)# qos enable
FC_switch_B_1(config)# qos class-map FCVI_1_10_Class match-any
FC_switch_B_1(config)# qos policy-map FCVI_1_10_Policy
FC_switch_B_1(config-pmap)# class FCVI_1_10_Class
FC_switch_B_1(config-pmap-c)# priority high
FC_switch_B_1(config-pmap-c)# exit
FC_switch_B_1(config)# exit
FC_switch_B_1(config)# qos service policy FCVI_1_10_Policy vsan 10
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
```

6. Configure the storage VSAN:

a. Set the VSAN ID:

```
vsan vsan-ID
```

b. Set the VSAN name:

```
vsan vsan-ID name vsan_name
```

The commands to configure the storage VSAN on FC_switch_A_1:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 20
FC_switch_A_1(config-vsan-db)# vsan 20 name STOR_1_20
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to configure the storage VSAN on FC_switch_B_1:

```
FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 20
FC_switch_B_1(config-vsan-db)# vsan 20 name STOR_1_20
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config
```

7. Add ports to the storage VSAN.

For the storage VSAN, all ports connecting HBA or FC-to-SAS bridges must be added. In this example fc1/5, fc1/9, fc1/17, fc1/21, fc1/25, fc1/29, fc1/33, and fc1/37 are being added.

The commands to add ports to the storage VSAN on FC_switch_A_1:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config)# vsan 20 interface fc1/5
FC_switch_A_1(config)# vsan 20 interface fc1/9
FC_switch_A_1(config)# vsan 20 interface fc1/17
FC_switch_A_1(config)# vsan 20 interface fc1/21
FC_switch_A_1(config)# vsan 20 interface fc1/25
FC_switch_A_1(config)# vsan 20 interface fc1/29
FC_switch_A_1(config)# vsan 20 interface fc1/33
FC_switch_A_1(config)# vsan 20 interface fc1/37
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to add ports to the storage VSAN on FC_switch_B_1:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config)# vsan 20 interface fc1/5
FC_switch_B_1(config)# vsan 20 interface fc1/9
FC_switch_B_1(config)# vsan 20 interface fc1/17
FC_switch_B_1(config)# vsan 20 interface fc1/21
FC_switch_B_1(config)# vsan 20 interface fc1/25
FC_switch_B_1(config)# vsan 20 interface fc1/29
FC_switch_B_1(config)# vsan 20 interface fc1/33
FC_switch_B_1(config)# vsan 20 interface fc1/37
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config

```

Configuring E-ports

You must configure the switch ports that connect the ISL (these are the E-Ports).

About this task

The procedure you use depends on which switch you are using:

- [Configuring the E-ports on the Cisco FC switch](#)
- [Configuring FCIP ports for a single ISL on Cisco 9250i FC switches](#)
- [Configuring FCIP ports for a dual ISL on Cisco 9250i FC switches](#)

Configuring the E-ports on the Cisco FC switch

You must configure the FC switch ports that connect the inter-switch link (ISL).

About this task

These are the E-ports, and configuration must be done for each port. To do so, you must calculate the correct number of buffer-to-buffer credits (BBCs).

All ISLs in the fabric must be configured with the same speed and distance settings.

This task must be performed on each ISL port.

Steps

1. Use the following table to determine the adjusted required BBCs per kilometer for possible port speeds.

To determine the correct number of BBCs, you multiply the Adjusted BBCs required (determined from the following table) by the distance in kilometers between the switches. An adjustment factor of 1.5 is required to account for FC-VI framing behavior.

Speed in Gbps	BBCs required per kilometer	Adjusted BBCs required (BBCs per km x 1.5)
1	0.5	0.75

2	1	1.5
4	2	3
8	4	6
16	8	12

For example, to compute the required number of credits for a distance of 30 km on a 4-Gbps link, make the following calculation:

- Speed in Gbps is 4
- Adjusted BBCs required is 3
- Distance in kilometers between switches is 30 km
- $3 \times 30 = 90$

1. Enter configuration mode:

```
config t
```

2. Specify the port you are configuring:

```
interface port-name
```

3. Shut down the port:

```
shutdown
```

4. Set the rate mode of the port to "dedicated":

```
switchport rate-mode dedicated
```

5. Set the speed for the port:

```
switchport speed speed-value
```

6. Set the buffer-to-buffer credits for the port:

```
switchport fcrxbbcredit number_of_buffers
```

7. Set the port to E mode:

```
switchport mode E
```

8. Enable the trunk mode for the port:

```
switchport trunk mode on
```

9. Add the ISL virtual storage area networks (VSANs) to the trunk:

```
switchport trunk allowed vsan 10
```

```
switchport trunk allowed vsan add 20
```

10. Add the port to port channel 1:

```
channel-group 1
```

11. Repeat the previous steps for the matching ISL port on the partner switch in the fabric.

The following example shows port fc1/41 configured for a distance of 30 km and 8 Gbps:

```
FC_switch_A_1# conf t
FC_switch_A_1# shutdown
FC_switch_A_1# switchport rate-mode dedicated
FC_switch_A_1# switchport speed 8000
FC_switch_A_1# switchport fcrxbbcredit 60
FC_switch_A_1# switchport mode E
FC_switch_A_1# switchport trunk mode on
FC_switch_A_1# switchport trunk allowed vsan 10
FC_switch_A_1# switchport trunk allowed vsan add 20
FC_switch_A_1# channel-group 1
fc1/36 added to port-channel 1 and disabled

FC_switch_B_1# conf t
FC_switch_B_1# shutdown
FC_switch_B_1# switchport rate-mode dedicated
FC_switch_B_1# switchport speed 8000
FC_switch_B_1# switchport fcrxbbcredit 60
FC_switch_B_1# switchport mode E
FC_switch_B_1# switchport trunk mode on
FC_switch_B_1# switchport trunk allowed vsan 10
FC_switch_B_1# switchport trunk allowed vsan add 20
FC_switch_B_1# channel-group 1
fc1/36 added to port-channel 1 and disabled
```

12. Issue the following command on both switches to restart the ports:

```
no shutdown
```

13. Repeat the previous steps for the other ISL ports in the fabric.
14. Add the native VSAN to the port-channel interface on both switches in the same fabric:

```
interface port-channel number
```

```
switchport trunk allowed vsan add native_san_id
```

15. Verify configuration of the port-channel:

```
show interface port-channel number
```

The port channel should have the following attributes:

- The port-channel is "trunking".
- Admin port mode is E, trunk mode is on.
- Speed shows the cumulative value of all the ISL link speeds.

For example, two ISL ports operating at 4 Gbps should show a speed of 8 Gbps.

- Trunk vsans (admin allowed and active) shows all the allowed VSANs.
- Trunk vsans (up) shows all the allowed VSANs.
- The member list shows all the ISL ports that were added to the port-channel.
- The port VSAN number should be the same as the VSAN that contains the ISLs (usually native vsan 1).

```
FC_switch_A_1(config-if)# show int port-channel 1
port-channel 1 is trunking
  Hardware is Fibre Channel
  Port WWN is 24:01:54:7f:ee:e2:8d:a0
  Admin port mode is E, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Speed is 8 Gbps
  Trunk vsans (admin allowed and active) (1,10,20)
  Trunk vsans (up) (1,10,20)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  5 minutes input rate 1154832 bits/sec,144354 bytes/sec, 170
frames/sec
  5 minutes output rate 1299152 bits/sec,162394 bytes/sec, 183
frames/sec
  535724861 frames input,1069616011292 bytes
    0 discards,0 errors
    0 invalid CRC/FCS,0 unknown class
    0 too long,0 too short
  572290295 frames output,1144869385204 bytes
    0 discards,0 errors
  5 input OLS,11 LRR,2 NOS,0 loop inits
  14 output OLS,5 LRR, 0 NOS, 0 loop inits
  Member[1] : fc1/36
  Member[2] : fc1/40
  Interface last changed at Thu Oct 16 11:48:00 2014
```

1. Exit interface configuration on both switches:

```
end
```

2. Copy the updated configuration to the startup configuration on both fabrics:

```
copy running-config startup-config
```

```
FC_switch_A_1(config-if)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1(config-if)# end
FC_switch_B_1# copy running-config startup-config
```

3. Repeat the previous steps on the second switch fabric.

Related information

You need to verify that you are using the specified port assignments when you cable the FC switches. Refer to [Port assignments for FC switches](#)

Configuring FCIP ports for a single ISL on Cisco 9250i FC switches

You must configure the FCIP switch ports that connect the ISL (E-ports) by creating FCIP profiles and interfaces, and then assigning them to the IPStorage1/1 GbE interface.

About this task

This task is only for configurations using a single ISL per switch fabric, using the IPStorage1/1 interface on each switch.

This task must be performed on each FC switch.

Two FCIP profiles are created on each switch:

- Fabric 1
 - FC_switch_A_1 is configured with FCIP profiles 11 and 111.
 - FC_switch_B_1 is configured with FCIP profiles 12 and 121.
- Fabric 2
 - FC_switch_A_2 is configured with FCIP profiles 13 and 131.
 - FC_switch_B_2 is configured with FCIP profiles 14 and 141.

Steps

1. Enter configuration mode:

```
config t
```

2. Enable FCIP:

```
feature fcip
```

3. Configure the IPStorage1/1 GbE interface:

- a. Enter configuration mode:

```
conf t
```

- b. Specify the IPStorage1/1 interface:

```
interface IPStorage1/1
```

- c. Specify the IP address and subnet mask:

```
interface ip-address subnet-mask
```

- d. Specify the MTU size of 2500:

```
switchport mtu 2500
```

- e. Enable the port:

```
no shutdown
```

- f. Exit configuration mode:

```
exit
```

The following example shows the configuration of an IPStorage1/1 port:

```
conf t
interface IPStorage1/1
  ip address 192.168.1.201 255.255.255.0
  switchport mtu 2500
  no shutdown
exit
```

4. Configure the FCIP profile for FC-VI traffic:

- a. Configure an FCIP profile and enter FCIP profile configuration mode:

```
fcip profile FCIP-profile-name
```

The profile name depends on which switch is being configured.

- b. Assign the IP address of the IPStorage1/1 interface to the FCIP profile:

```
ip address ip-address
```

- c. Assign the FCIP profile to TCP port 3227:

```
port 3227
```

- d. Set the TCP settings:

```

tcp keepalive-timeout 1

tcp max-retransmissions 3

max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-ms
3

tcp min-retransmit-time 200

tcp keepalive-timeout 1

tcp pmtu-enable reset-timeout 3600

tcp sack-enable ``no tcp cwm

```

The following example shows the configuration of the FCIP profile:

```

conf t
fcip profile 11
  ip address 192.168.1.333
  port 3227
  tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-
time-ms 3
  tcp min-retransmit-time 200
  tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
no tcp cwm

```

5. Configure the FCIP profile for storage traffic:

- a. Configure an FCIP profile with the name 111 and enter FCIP profile configuration mode:

```
fcip profile 111
```

- b. Assign the IP address of the IPStorage1/1 interface to the FCIP profile:

```
ip address ip-address
```

- c. Assign the FCIP profile to TCP port 3229:

```
port 3229
```

- d. Set the TCP settings:

```
tcp keepalive-timeout 1
```

```

tcp max-retransmissions 3

max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-ms
3

tcp min-retransmit-time 200

tcp keepalive-timeout 1

tcp pmtu-enable reset-timeout 3600

tcp sack-enable ``no tcp cwm

```

The following example shows the configuration of the FCIP profile:

```

conf t
fcip profile 111
  ip address 192.168.1.334
  port 3229
  tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-
time-ms 3
  tcp min-retransmit-time 200
  tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
  no tcp cwm

```

6. Create the first of two FCIP interfaces:

```
interface fcip 1
```

This interface is used for FC-IV traffic.

a. Select the profile 11 created previously:

```
use-profile 11
```

b. Set the IP address and port of the IPStorage1/1 port on the partner switch:

```
peer-info ipaddr partner-switch-port-ip port 3227
```

c. Select TCP connection 2:

```
tcp-connection 2
```

d. Disable compression:

```
no ip-compression
```

e. Enable the interface:

```
no shutdown
```

f. Configure the control TCP connection to 48 and the data connection to 26 to mark all packets on that differentiated services code point (DSCP) value:

```
qos control 48 data 26
```

g. Exit the interface configuration mode:

```
exit
```

The following example shows the configuration of the FCIP interface:

```
interface fcip 1
  use-profile 11
  # the port # listed in this command is the port that the remote switch
  is listening on
  peer-info ipaddr 192.168.32.334   port 3227
  tcp-connection 2
  no ip-compression
  no shutdown
  qos control 48 data 26
exit
```

7. Create the second of two FCIP interfaces:

```
interface fcip 2
```

This interface is used for storage traffic.

a. Select the profile 111 created previously:

```
use-profile 111
```

b. Set the IP address and port of the IPStorage1/1 port on the partner switch:

```
peer-info ipaddr partner-switch-port-ip port 3229
```

c. Select TCP connection 2:

```
tcp-connection 5
```

d. Disable compression:

```
no ip-compression
```

e. Enable the interface:

```
no shutdown
```

- f. Configure the control TCP connection to 48 and data connection to 26 to mark all packets on that differentiated services code point (DSCP) value:

```
qos control 48 data 26
```

- g. Exit the interface configuration mode:

```
exit
```

The following example shows the configuration of the FCIP interface:

```
interface fcip 2
  use-profile 11
  # the port # listed in this command is the port that the remote switch
  is listening on
  peer-info ipaddr 192.168.32.33e port 3229
  tcp-connection 5
  no ip-compression
  no shutdown
  qos control 48 data 26
exit
```

8. Configure the switchport settings on the fcip 1 interface:

- a. Enter configuration mode:

```
config t
```

- b. Specify the port you are configuring:

```
interface fcip 1
```

- c. Shut down the port:

```
shutdown
```

- d. Set the port to E mode:

```
switchport mode E
```

- e. Enable the trunk mode for the port:

```
switchport trunk mode on
```

- f. Set the trunk allowed vsan to 10:

```
switchport trunk allowed vsan 10
```

- g. Set the speed for the port:

```
switchport speed speed-value
```

9. Configure the switchport settings on the fcip 2 interface:

a. Enter configuration mode:

```
config t
```

b. Specify the port you are configuring:

```
interface fcip 2
```

c. Shut down the port:

```
shutdown
```

d. Set the port to E mode:

```
switchport mode E
```

e. Enable the trunk mode for the port:

```
switchport trunk mode on
```

f. Set the trunk allowed vsan to 20:

```
switchport trunk allowed vsan 20
```

g. Set the speed for the port:

```
switchport speed speed-value
```

10. Repeat the previous steps on the second switch.

The only differences are the appropriate IP addresses and unique FCIP profile names.

- When configuring the first switch fabric, FC_switch_B_1 is configured with FCIP profiles 12 and 121.
- When configuring the first switch fabric, FC_switch_A_2 is configured with FCIP profiles 13 and 131 and FC_switch_B_2 is configured with FCIP profiles 14 and 141.

11. Restart the ports on both switches:

```
no shutdown
```

12. Exit the interface configuration on both switches:

```
end
```

13. Copy the updated configuration to the startup configuration on both switches:

```
copy running-config startup-config
```

```

FC_switch_A_1(config-if)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1(config-if)# end
FC_switch_B_1# copy running-config startup-config

```

14. Repeat the previous steps on the second switch fabric.

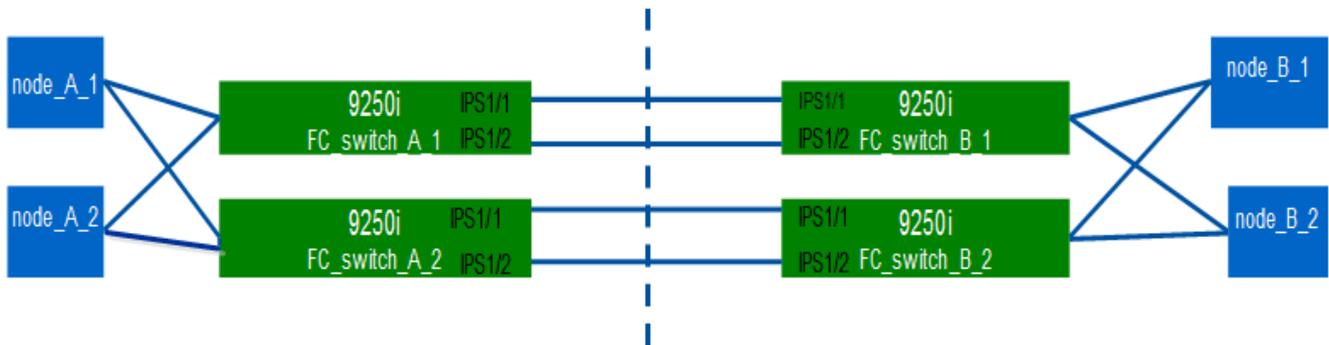
Configuring FCIP ports for a dual ISL on Cisco 9250i FC switches

You must configure the FCIP switch ports that connect the ISL (E-ports) by creating FCIP profiles and interfaces, and then assigning them to the IPStorage1/1 and IPStorage1/2 GbE interfaces.

About this task

This task is only for configurations that use a dual ISL per switch fabric, using the IPStorage1/1 and IPStorage1/2 GbE interfaces on each switch.

This task must be performed on each FC switch.



The task and examples use the following profile configuration tables:

- [Fabric 1 profile configuration table](#)
- [Fabric 2 profile configuration table](#)

Fabric 1 profile configuration table

Switch fabric	IPStorage interface	IP Address	Port type	FCIP interface	FCIP profile	Port	Peer IP/port	VSAN ID
---------------	---------------------	------------	-----------	----------------	--------------	------	--------------	---------

FC_switch_A_1	IPStorage 1/1	a.a.a.a	FC-VI	fcip 1	15	3220	c.c.c.c/3230	10
			Storage	fcip 2	20	3221	c.c.c.c/3231	20
	IPStorage 1/2	b.b.b.b	FC-VI	fcip 3	25	3222	d.d.d.d/3232	10
			Storage	fcip 4	30	3223	d.d.d.d/3233	20
FC_switch_B_1	IPStorage 1/1	c.c.c.c	FC-VI	fcip 1	15	3230	a.a.a.a/3220	10
			Storage	fcip 2	20	3231	a.a.a.a/3221	20
	IPStorage 1/2	d.d.d.d	FC-VI	fcip 3	25	3232	b.b.b.b/3222	10
			Storage	fcip 4	30	3233	b.b.b.b/3223	20

Fabric 2 profile configuration table

Switch fabric	IPStorage interface	IP Address	Port type	FCIP interface	FCIP profile	Port	Peer IP/port	VSAN ID
FC_switch_A_2	IPStorage 1/1	e.e.e.e	FC-VI	fcip 1	15	3220	g.g.g.g/3230	10
			Storage	fcip 2	20	3221	g.g.g.g/3231	20
	IPStorage 1/2	f.f.f.f	FC-VI	fcip 3	25	3222	h.h.h.h/3232	10
			Storage	fcip 4	30	3223	h.h.h.h/3233	20

FC_switch _B_2	IPStorage 1/1	g.g.g.g	FC-VI	fcip 1	15	3230	e.e.e.e/32 20	10
			Storage	fcip 2	20	3231	e.e.e.e/32 21	20
	IPStorage 1/2	h.h.h.h	FC-VI	fcip 3	25	3232	f.f.f.f/3222	10
			Storage	fcip 4	30	3233	f.f.f.f/3223	20

Steps

1. Enter configuration mode:

```
config t
```

2. Enable FCIP:

```
feature fcip
```

3. On each switch, configure the two IPStorage interfaces ("IPStorage1/1" and "IPStorage1/2"):

- a. Enter configuration mode:

```
conf t
```

- b. Specify the IPStorage interface to create:

```
interface ipstorage
```

The *ipstorage* parameter value is "IPStorage1/1" or "IPStorage1/2".

- c. Specify the IP address and subnet mask of the IPStorage interface previously specified:

```
interface ip-address subnet-mask
```



On each switch, the IPStorage interfaces "IPStorage1/1" and "IPStorage1/2" must have different IP addresses.

- d. Specify the MTU size as 2500:

```
switchport mtu 2500
```

- e. Enable the port:

```
no shutdown
```

- f. Exit configuration mode:

```
exit
```

- g. Repeat [Substep "a"](#) through [Substep "f"](#) to configure the IPStorage1/2 GbE interface with a different IP

address.

4. Configure the FCIP profiles for FC-VI and storage traffic with the profile names given in the profile configuration table:

- a. Enter configuration mode:

```
conf t
```

- b. Configure the FCIP profiles with the following profile names:

```
fcip profile FCIP-profile-name
```

The following list provides the values for the *FCIP-profile-name* parameter:

- 15 for FC-VI on IPStorage1/1
- 20 for storage on IPStorage1/1
- 25 for FC-VI on IPStorage1/2
- 30 for storage on IPStorage1/2

- c. Assign the FCIP profile ports according to the profile configuration table:

```
port port_number
```

- d. Set the TCP settings:

```
tcp keepalive-timeout 1
```

```
tcp max-retransmissions 3
```

```
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-ms 3
```

```
tcp min-retransmit-time 200
```

```
tcp keepalive-timeout 1
```

```
tcp pmtu-enable reset-timeout 3600
```

```
tcp sack-enable
```

```
no tcp cwm
```

5. Create FCIP interfaces:

```
interface fcip FCIP_interface
```

The *FCIP_interface* parameter value is “1”, “2”, “3”, or “4” as shown in the profile configuration table.

- a. Map interfaces to the previously created profiles:

```
use-profile profile
```

- b. Set the peer IP address and peer profile port number:

```
peer-info peer IPstorage ipaddr port peer_profile_port_number
```

- c. Select the TCP connections:

```
tcp-connection connection-#
```

The *connection-#* parameter value is “2” for FC-VI profiles and “5” for storage profiles.

- d. Disable compression:

```
no ip-compression
```

- e. Enable the interface:

```
no shutdown
```

- f. Configure the control TCP connection to “48” and the data connection to “26” to mark all packets that have differentiated services code point (DSCP) value:

```
qos control 48 data 26
```

- g. Exit configuration mode:

```
exit
```

6. Configure the switchport settings on each FCIP interface:

- a. Enter configuration mode:

```
config t
```

- b. Specify the port that you are configuring:

```
interface fcip 1
```

- c. Shut down the port:

```
shutdown
```

- d. Set the port to E mode:

```
switchport mode E
```

- e. Enable the trunk mode for the port:

```
switchport trunk mode on
```

- f. Specify the trunk that is allowed on a specific VSAN:

```
switchport trunk allowed vsan vsan_id
```

The *vsan_id* parameter value is “VSAN 10” for FC-VI profiles and “VSAN 20” for storage profiles.

- g. Set the speed for the port:

```
switchport speed speed-value
```

h. Exit configuration mode:

```
exit
```

7. Copy the updated configuration to the startup configuration on both switches:

```
copy running-config startup-config
```

The following examples show the configuration of FCIP ports for a dual ISL in fabric 1 switches FC_switch_A_1 and FC_switch_B_1.

For FC_switch_A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# no in-order-guarantee vsan 10
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config

# fcip settings

feature fcip

conf t
interface IPStorage1/1
# IP address: a.a.a.a
# Mask: y.y.y.y
ip address <a.a.a.a y.y.y.y>
switchport mtu 2500
no shutdown
exit
conf t
fcip profile 15
ip address <a.a.a.a>
port 3220
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable
no tcp cwm

conf t
fcip profile 20
```

```

ip address <a.a.a.a>
port 3221
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable
no tcp cwm

conf t
interface IPStorage1/2
# IP address: b.b.b.b
# Mask: y.y.y.y
ip address <b.b.b.b y.y.y.y>
switchport mtu 2500
no shutdown
exit

conf t
fcip profile 25
ip address <b.b.b.b>
port 3222
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable
no tcp cwm

conf t
fcip profile 30
ip address <b.b.b.b>
port 3223
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600

```

```
    tcp sack-enable
    no tcp cwm
interface fcip 1
    use-profile 15
# the port # listed in this command is the port that the remote switch is
listening on
    peer-info ipaddr <c.c.c.c> port 3230
    tcp-connection 2
    no ip-compression
    no shutdown
    qos control 48 data 26
exit

interface fcip 2
    use-profile 20
# the port # listed in this command is the port that the remote switch is
listening on
    peer-info ipaddr <c.c.c.c> port 3231
    tcp-connection 5
    no ip-compression
    no shutdown
    qos control 48 data 26
exit

interface fcip 3
    use-profile 25
# the port # listed in this command is the port that the remote switch is
listening on
    peer-info ipaddr < d.d.d.d > port 3232
    tcp-connection 2
    no ip-compression
    no shutdown
    qos control 48 data 26
exit

interface fcip 4
    use-profile 30
# the port # listed in this command is the port that the remote switch is
listening on
    peer-info ipaddr < d.d.d.d > port 3233
    tcp-connection 5
    no ip-compression
    no shutdown
    qos control 48 data 26
exit
```

```
conf t
interface fcip 1
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit
```

```
conf t
interface fcip 2
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit
```

```
conf t
interface fcip 3
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit
```

```
conf t
interface fcip 4
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit
```

For FC_switch_B_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# in-order-guarantee vsan 10
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config

# fcip settings
```

```

feature fcip

conf t
interface IPStorage1/1
# IP address: c.c.c.c
# Mask: y.y.y.y
ip address <c.c.c.c y.y.y.y>
switchport mtu 2500
no shutdown
exit

conf t
fcip profile 15
ip address <c.c.c.c>
port 3230
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable
no tcp cwm

conf t
fcip profile 20
ip address <c.c.c.c>
port 3231
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
tcp min-retransmit-time 200
tcp keepalive-timeout 1
tcp pmtu-enable reset-timeout 3600
tcp sack-enable
no tcp cwm

conf t
interface IPStorage1/2
# IP address: d.d.d.d
# Mask: y.y.y.y
ip address <b.b.b.b y.y.y.y>
switchport mtu 2500
no shutdown

```

```

exit

conf t
fcip profile 25
  ip address <d.d.d.d>
  port 3232
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
  tcp min-retransmit-time 200
  tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
  no tcp cwm

conf t
fcip profile 30
  ip address <d.d.d.d>
  port 3233
tcp keepalive-timeout 1
tcp max-retransmissions 3
max-bandwidth-mbps 5000 min-available-bandwidth-mbps 4500 round-trip-time-
ms 3
  tcp min-retransmit-time 200
  tcp keepalive-timeout 1
  tcp pmtu-enable reset-timeout 3600
  tcp sack-enable
  no tcp cwm

interface fcip 1
  use-profile 15
# the port # listed in this command is the port that the remote switch is
listening on
  peer-info ipaddr <a.a.a.a> port 3220
  tcp-connection 2
  no ip-compression
  no shutdown
  qos control 48 data 26
exit

interface fcip 2
  use-profile 20
# the port # listed in this command is the port that the remote switch is
listening on
  peer-info ipaddr <a.a.a.a> port 3221

```

```

    tcp-connection 5
    no ip-compression
    no shutdown
    qos control 48 data 26
exit

interface fcip 3
    use-profile 25
# the port # listed in this command is the port that the remote switch is
listening on
    peer-info ipaddr < b.b.b.b > port 3222
    tcp-connection 2
    no ip-compression
    no shutdown
    qos control 48 data 26
exit

interface fcip 4
    use-profile 30
# the port # listed in this command is the port that the remote switch is
listening on
    peer-info ipaddr < b.b.b.b > port 3223
    tcp-connection 5
    no ip-compression
    no shutdown
    qos control 48 data 26
exit

conf t
interface fcip 1
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit

conf t
interface fcip 2
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit

```

```

conf t
interface fcip 3
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 10
no shutdown
exit

conf t
interface fcip 4
shutdown
switchport mode E
switchport trunk mode on
switchport trunk allowed vsan 20
no shutdown
exit

```

Configuring zoning on a Cisco FC switch

You must assign the switch ports to separate zones to isolate storage (HBA) and controller (FC-VI) traffic.

About this task

These steps must be performed on both FC switch fabrics.

The following steps use the zoning described in the section Zoning for a FibreBridge 7500N in a four-node MetroCluster configuration.

Steps

1. Clear the existing zones and zone set, if present.
 - a. Determine which zones and zone sets are active:

```
show zoneset active
```

```
FC_switch_A_1# show zoneset active
```

```
FC_switch_B_1# show zoneset active
```

- b. Disable the active zone sets identified in the previous step:

```
no zoneset activate name zoneset_name vsan vsan_id
```

The following example shows two zone sets being disabled:

- ZoneSet_A on FC_switch_A_1 in VSAN 10
- ZoneSet_B on FC_switch_B_1 in VSAN 20

```
FC_switch_A_1# no zoneset activate name ZoneSet_A vsan 10

FC_switch_B_1# no zoneset activate name ZoneSet_B vsan 20
```

c. After all zone sets are deactivated, clear the zone database:

```
clear zone database zone-name
```

```
FC_switch_A_1# clear zone database 10
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# clear zone database 20
FC_switch_B_1# copy running-config startup-config
```

2. Obtain the switch worldwide name (WWN):

```
show wwn switch
```

3. Configure the basic zone settings:

a. Set the default zoning policy to “permit”:

```
no system default zone default-zone permit
```

b. Enable the full zone distribution:

```
system default zone distribute full
```

c. Set the default zoning policy for each VSAN:

```
no zone default-zone permit vsanid
```

d. Set the default full zone distribution for each VSAN:

```
zoneset distribute full vsanid
```

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# no system default zone default-zone permit
FC_switch_A_1(config)# system default zone distribute full
FC_switch_A_1(config)# no zone default-zone permit 10
FC_switch_A_1(config)# no zone default-zone permit 20
FC_switch_A_1(config)# zoneset distribute full vsan 10
FC_switch_A_1(config)# zoneset distribute full vsan 20
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config

FC_switch_B_1# conf t
FC_switch_B_1(config)# no system default zone default-zone permit
FC_switch_B_1(config)# system default zone distribute full
FC_switch_B_1(config)# no zone default-zone permit 10
FC_switch_B_1(config)# no zone default-zone permit 20
FC_switch_B_1(config)# zoneset distribute full vsan 10
FC_switch_B_1(config)# zoneset distribute full vsan 20
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config

```

4. Create storage zones and add the storage ports to them.



Perform these steps on only one switch in each fabric.

The zoning depends on the model FC-to-SAS bridge you are using. For details, see the section for your model bridge. The examples show Brocade switch ports, so adjust your ports accordingly.

- [Zoning for FibreBridge 7500N or 7600N bridges using one FC port](#)
- [Zoning for FibreBridge 7500N bridges using both FC ports](#)

Each storage zone contains the HBA initiator ports from all controllers and one single port connecting an FC-to-SAS bridge.

a. Create the storage zones:

```
zone name STOR-zone-name vsan vsanid
```

b. Add storage ports to the zone:

```
member portswitch WWN
```

c. Activate the zone set:

```
zoneset activate name STOR-zone-name-setname vsan vsan-id
```

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# zone name STOR_Zone_1_20_25 vsan 20
FC_switch_A_1(config-zone)# member interface fc1/5 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/9 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/17 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/21 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/5 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/9 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/17 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/21 swwn
20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/25 swwn
20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# end
FC_switch_A_1# copy running-config startup-config

```

5. Create a storage zone set and add the storage zones to the new set.



Perform these steps on only one switch in the fabric.

a. Create the storage zone set:

```
zoneset name STOR-zone-set-name vsan vsan-id
```

b. Add storage zones to the zone set:

```
member STOR-zone-name
```

c. Activate the zone set:

```
zoneset activate name STOR-zone-set-name vsan vsanid
```

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# zoneset name STORI_Zoneset_1_20 vsan 20
FC_switch_A_1(config-zoneset)# member STOR_Zone_1_20_25
...
FC_switch_A_1(config-zoneset)# exit
FC_switch_A_1(config)# zoneset activate name STOR_ZoneSet_1_20 vsan
20
FC_switch_A_1(config)# exit
FC_switch_A_1# copy running-config startup-config

```

6. Create FCVI zones and add the FCVI ports to them.

Each FCVI zone contains the FCVI ports from all the controllers of one DR Group.



Perform these steps on only one switch in the fabric.

The zoning depends on the model FC-to-SAS bridge you are using. For details, see the section for your model bridge. The examples show Brocade switch ports, so adjust your ports accordingly.

- [Zoning for FibreBridge 7500N or 7600N bridges using one FC port](#)
- [Zoning for FibreBridge 7500N bridges using both FC ports](#)

Each storage zone contains the HBA initiator ports from all controllers and one single port connecting an FC-to-SAS bridge.

a. Create the FCVI zones:

```
zone name FCVI-zone-name vsan vsanid
```

b. Add FCVI ports to the zone:

```
member FCVI-zone-name
```

c. Activate the zone set:

```
zoneset activate name FCVI-zone-name-set-name vsan vsanid
```

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# zone name FCVI_Zone_1_10_25 vsan 10
FC_switch_A_1(config-zone)# member interface fc1/1
swwn20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/2
swwn20:00:00:05:9b:24:cb:78
FC_switch_A_1(config-zone)# member interface fc1/1
swwn20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# member interface fc1/2
swwn20:00:00:05:9b:24:12:99
FC_switch_A_1(config-zone)# end
FC_switch_A_1# copy running-config startup-config

```

7. Create an FCVI zone set and add the FCVI zones to it:



Perform these steps on only one switch in the fabric.

a. Create the FCVI zone set:

```
zoneset name FCVI_zone_set_name vsan vsan-id
```

b. Add FCVI zones to the zone set:

```
member FCVI_zonename
```

c. Activate the zone set:

```
zoneset activate name FCVI_zone_set_name vsan vsan-id
```

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# zoneset name FCVI_Zoneset_1_10 vsan 10
FC_switch_A_1(config-zoneset)# member FCVI_Zone_1_10_25
FC_switch_A_1(config-zoneset)# member FCVI_Zone_1_10_29
...
FC_switch_A_1(config-zoneset)# exit
FC_switch_A_1(config)# zoneset activate name FCVI_ZoneSet_1_10 vsan 10
FC_switch_A_1(config)# exit
FC_switch_A_1# copy running-config startup-config

```

8. Verify the zoning:

```
show zone
```

9. Repeat the previous steps on the second FC switch fabric.

Ensuring the FC switch configuration is saved

You must make sure the FC switch configuration is saved to the startup config on all switches.

Step

Issue the following command on both FC switch fabrics:

```
copy running-config startup-config
```

```
FC_switch_A_1# copy running-config startup-config
```

```
FC_switch_B_1# copy running-config startup-config
```

Install FC-to-SAS bridges and SAS disk shelves

You install and cable ATTO FibreBridge bridges and SAS disk shelves when adding new storage to the configuration.

About this task

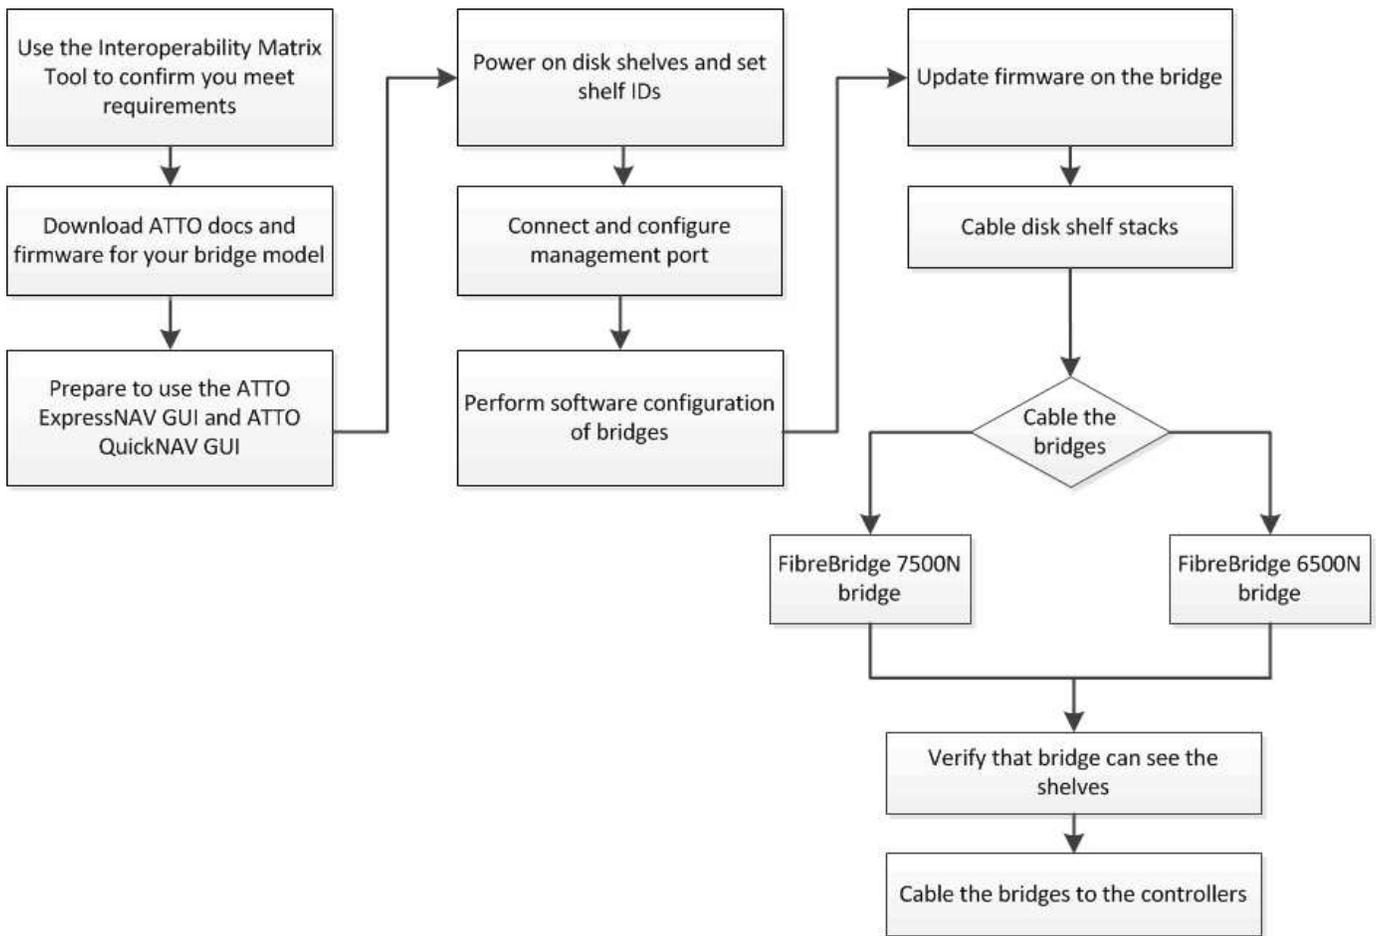
For systems received from the factory, the FC-to-SAS bridges are preconfigured and do not require additional configuration.

This procedure is written with the assumption that you are using the recommended bridge management interfaces: the ATTO ExpressNAV GUI and ATTO QuickNAV utility.

You use the ATTO ExpressNAV GUI to configure and manage a bridge, and to update the bridge firmware. You use the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port.

You can use other management interfaces instead, if needed, such as a serial port or Telnet to configure and manage a bridge and to configure the Ethernet management 1 port, and FTP to update the bridge firmware.

This procedure uses the following workflow:



In-band management of the FC-to-SAS bridges

Beginning with ONTAP 9.5 with FibreBridge 7500N or 7600N bridges, *in-band management* of the bridges is supported as an alternative to IP management of the bridges. Beginning with ONTAP 9.8, out-of-band management is deprecated.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

When using in-band management, the bridges can be managed and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required, reducing the security vulnerability of the bridge.

The availability of in-band management of the bridges depends on the version of ONTAP:

- Beginning with ONTAP 9.8, bridges are managed via in-band connections by default and out-of-band management of the bridges via SNMP is deprecated.
- ONTAP 9.5 through 9.7: Either in-band management or out-of-band SNMP management is supported.
- Prior to ONTAP 9.5, only out-of-band SNMP management is supported.

Bridge CLI commands can be issued from the ONTAP interface `storage bridge run-cli -name bridge_name -command bridge_command_name command` at the ONTAP interface.



Using in-band management with IP access disabled is recommended to improve security by limiting physical connectivity to the bridge.

FibreBridge 7600N and 7500N bridge limits and attachment rules

Review the limits and considerations when attaching FibreBridge 7600N and 7500N bridges.

FibreBridge 7600N and 7500N bridge limits

- The maximum number of HDD and SSD drives combined is 240.
- The maximum number of SSD drives is 96.
- The maximum number of SSDs per SAS port is 48.
- The maximum number of shelves per SAS port is 10.

FibreBridge 7600N and 7500N bridge attachment rules

- Do not mix SSD and HDD drives on the same SAS port.
- Distribute the shelves evenly across the SAS ports.
- You shouldn't have DS460 shelves on the same SAS port as other shelf types (for example, DS212 or DS224 shelves).

Example configuration

The following shows an example configuration for connecting four DS224 shelves with SSD drives and six DS224 shelves with HDD drives:

SAS port	Shelves and drives
SAS port-A	2x DS224 shelves with SSD drives
SAS port-B	2x DS224 shelves with SSD drives
SAS port-C	3x DS224 shelves with HDD drives
SAS port-D	3x DS224 shelves with HDD drives

Prepare for the installation

When you are preparing to install the bridges as part of your new MetroCluster system, you must ensure that your system meets certain requirements, including meeting setup and configuration requirements for the bridges. Other requirements include downloading the necessary documents, the ATTO QuickNAV utility, and the bridge firmware.

Before you begin

- Your system must already be installed in a rack if it was not shipped in a system cabinet.
- Your configuration must be using supported hardware models and software versions.

In the [NetApp Interoperability Matrix Tool \(IMT\)](#), you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- Each FC switch must have one FC port available for one bridge to connect to it.
- You must have familiarized yourself with how to handle SAS cables and the considerations and best

practices for installing and cabling disk shelves.

The *Installation and Service Guide* for your disk shelf model describes the considerations and best practices.

- The computer you are using to set up the bridges must be running an ATTO-supported web browser to use the ATTO ExpressNAV GUI.

The *ATTO Product Release Notes* have an up-to-date list of supported web browsers. You can access this document from the ATTO web site as described in the following steps.

Steps

1. Download the *Installation and Service Guide* for your disk shelf model:

- a. Access the ATTO web site using the link provided for your FibreBridge model and download the manual and the QuickNAV utility.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this and other content on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

2. Gather the hardware and information needed to use the recommended bridge management interfaces, the ATTO ExpressNAV GUI, and the ATTO QuickNAV utility:

- a. Determine a non-default user name and password (for accessing the bridges).

You should change the default user name and password.

- b. If configuring for IP management of the bridges, you need the shielded Ethernet cable provided with the bridges (which connects from the bridge Ethernet management 1 port to your network).
- c. If configuring for IP management of the bridges, you need an IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge.
- d. Disable VPN clients on the computer you are using for setup.

Active VPN clients cause the QuickNAV scan for bridges to fail.

Install the FC-to-SAS bridge and SAS shelves

After ensuring that the system meets all of the requirements in “Preparing for the installation”, you can install your new system.

About this task

- The disk and shelf configuration at both sites should be identical.

If a non-mirrored aggregate is used, the disk and shelf configuration at each site might be different.



All disks in the disaster recovery group must use the same type of connection and be visible to all of the nodes within the disaster recovery group, regardless of the disks being used for mirrored or non-mirrored aggregate.

- The system connectivity requirements for maximum distances for disk shelves, FC switches, and backup

tape devices using 50-micron, multimode fiber-optic cables, also apply to FibreBridge bridges.

NetApp Hardware Universe



In-band ACP is supported without additional cabling in the following shelves and FibreBridge 7500N or 7600N bridge:

- IOM12 (DS460C) behind a 7500N or 7600N bridge with ONTAP 9.2 and later
- IOM12 (DS212C and DS224C) behind a 7500N or 7600N bridge with ONTAP 9.1 and later



SAS shelves in MetroCluster configurations do not support ACP cabling.

Enable IP port access on the FibreBridge 7600N bridge if necessary

If you are using an ONTAP version prior to 9.5, or otherwise plan to use out-of-band access to the FibreBridge 7600N bridge using telnet or other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV), you can enable the access services via the console port.

About this task

Unlike the ATTO FibreBridge 7500N bridges, the FibreBridge 7600N bridge is shipped with all IP port protocols and services disabled.

Beginning with ONTAP 9.5, *in-band management* of the bridges is supported. This means the bridges can be configured and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required and the bridge user interfaces are not required.

Beginning with ONTAP 9.8, *in-band management* of the bridges is supported by default and out-of-band SNMP management is deprecated.

This task is required if you are **not** using in-band management to manage the bridges. In this case, you need to configure the bridge via the Ethernet management port.

Steps

1. Access the bridge console interface by connecting a serial cable to the serial port on the FibreBridge 7600N bridge.
2. Using the console, enable the access services, and then save the configuration:

```
set closeport none  
  
saveconfiguration
```

The `set closeport none` command enables all access services on the bridge.

3. Disable a service, if desired, by issuing the `set closeport` command and repeating the command as necessary until all desired services are disabled:

```
set closeport service
```

The `set closeport` command disables a single service at a time.

The parameter *service* can be specified as one of the following:

- `expressnav`
- `ftp`
- `icmp`
- `quicknav`
- `snmp`
- `telnet`

You can check whether a specific protocol is enabled or disabled by using the `get closeport` command.

4. If you are enabling SNMP, you must also issue following command:

```
set SNMP enabled
```

SNMP is the only protocol that requires a separate enable command.

5. Save the configuration:

```
saveconfiguration
```

Configure the FC-to-SAS bridges

Before cabling your model of the FC-to-SAS bridges, you must configure the settings in the FibreBridge software.

Before you begin

You should decide whether you will be using in-band management of the bridges.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.

Steps

1. Configure the serial console port on the ATTO FibreBridge by setting the port speed to 115000 bauds:

```
get serialportbaudrate
SerialPortBaudRate = 115200

Ready.

set serialportbaudrate 115200

Ready. *
saveconfiguration
Restart is necessary....
Do you wish to restart (y/n) ? y
```

2. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

3. If configuring for IP management, connect the Ethernet management 1 port on each bridge to your network by using an Ethernet cable.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

The Ethernet management 1 port enables you to quickly download the bridge firmware (using ATTO ExpressNAV or FTP management interfaces) and to retrieve core files and extract logs.

4. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

5. Configure the bridge.

You should make note of the user name and password that you designate.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

To set the IP address without the QuickNAV utility, you need to have a serial connection to the FibreBridge.

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
```

```
set ethernetspeed mp1 1000
```

b. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridge_name
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

6. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the FC port of the controller module to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate <port-number> <port-speed>
```

- b. If you are configuring a FibreBridge 7500N bridge, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode <port-number> ptp
```

- c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable <port-number>
```

The following example shows the disabling of FC port 2:

```
FCPortDisable 2

Fibre Channel Port 2 has been disabled.
```

- d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used.

If only SAS port A is used, then SAS ports B, C, and D must be disabled. The following example shows the disabling of SAS port B. You must similarly disable SAS ports C and D:

```
SASPortDisable b

SAS Port B has been disabled.
```

7. Secure access to the bridge and save the bridge's configuration. Choose an option from below depending on the version of ONTAP your system is running.

ONTAP version	Steps
ONTAP 9.5 or later	<p>a. View the status of the bridges:</p> <pre data-bbox="561 233 878 260">storage bridge show</pre> <p>The output shows which bridge is not secured.</p> <p>b. Secure the bridge:</p> <pre data-bbox="561 436 760 464">securebridge</pre>
ONTAP 9.4 or earlier	<p>a. View the status of the bridges:</p> <pre data-bbox="561 604 878 632">storage bridge show</pre> <p>The output shows which bridge is not secured.</p> <p>b. Check the status of the unsecured bridge's ports:</p> <pre data-bbox="561 808 626 835">info</pre> <p>The output shows the status of Ethernet ports MP1 and MP2.</p> <p>c. If Ethernet port MP1 is enabled, run:</p> <pre data-bbox="561 1012 1045 1039">set EthernetPort mp1 disabled</pre> <p>If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.</p> <p>d. Save the bridge's configuration.</p> <p>You must run the following commands:</p> <pre data-bbox="561 1314 846 1341">SaveConfiguration</pre> <pre data-bbox="561 1388 812 1415">FirmwareRestart</pre> <p>You are prompted to restart the bridge.</p>

- After completing MetroCluster configuration, use the `flashimages` command to check your version of FibreBridge firmware and, if the bridges are not using the latest supported version, update the firmware on all bridges in the configuration.

Maintain MetroCluster Components

Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM12 modules

After configuring the bridge, you can start cabling your new system.

About this task

For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

Steps

1. Daisy-chain the disk shelves in each stack:
 - a. Beginning with the logical first shelf in the stack, connect IOM A port 3 to IOM A port 1 on the next shelf until each IOM A in the stack is connected.
 - b. Repeat the previous substep for IOM B.
 - c. Repeat the previous substeps for each stack.

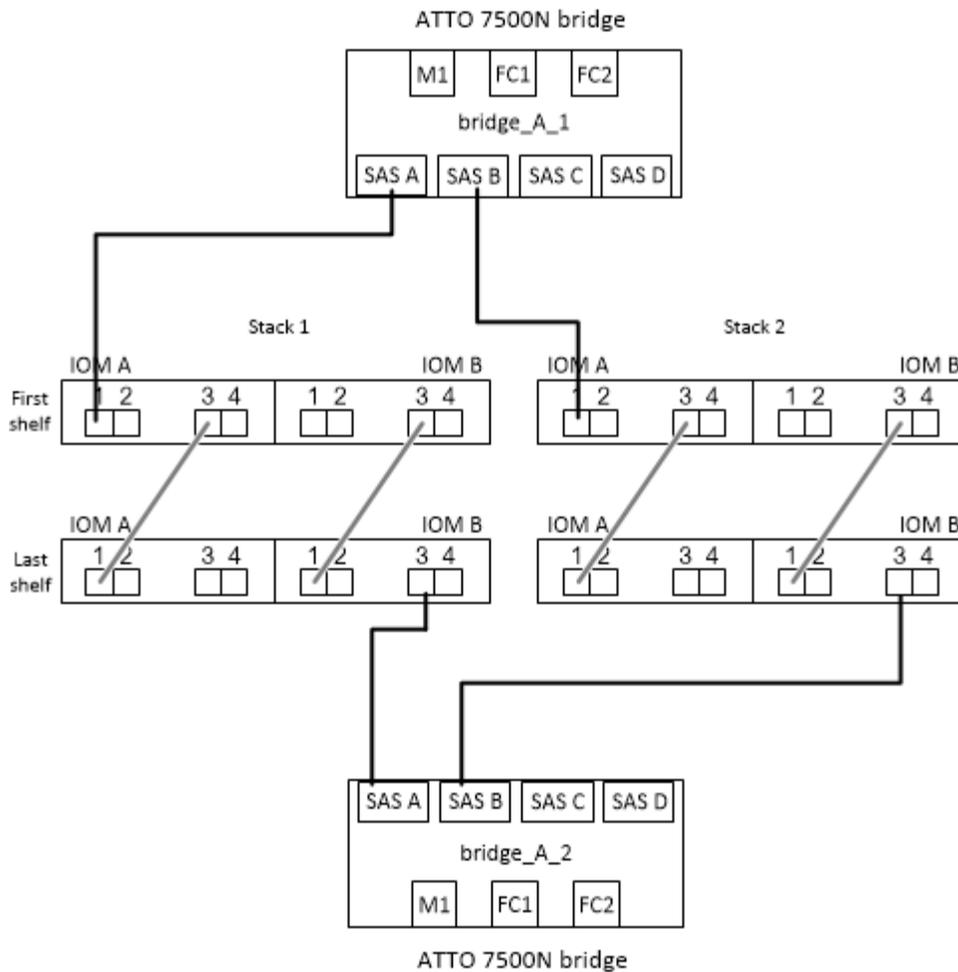
The *Installation and Service Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

2. Power on the disk shelves, and then set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites).
3. Cable disk shelves to the FibreBridge bridges.
 - a. For the first stack of disk shelves, cable IOM A of the first shelf to SAS port A on FibreBridge A, and cable IOM B of the last shelf to SAS port A on FibreBridge B.
 - b. For additional shelf stacks, repeat the previous step using the next available SAS port on the FibreBridge bridges, using port B for the second stack, port C for the third stack, and port D for the fourth stack.
 - c. During cabling, attach the stacks based on IOM12 modules to the same bridge as long as they are connected to separate SAS ports.



Each stack can use different models of IOM, but all disk shelves within a stack must use the same model.

The following illustration shows disk shelves connected to a pair of FibreBridge 7600N or 7500N bridges:



Verify bridge connectivity and cabling the bridge FC ports

You should verify that each bridge can detect all of the disk drives, and then cable each bridge to the local FC switches.

Steps

1. Verify that each bridge can detect all of the disk drives and disk shelves it is connected to:

If you are using the...	Then...
-------------------------	---------

ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of a bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge homepage of the bridge for which you entered the IP address, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page of the bridge appears with a menu to the left.</p> <p>c. Click Advanced.</p> <p>d. View the connected devices by using the <code>sastargets</code> command, and then click Submit.</p>
Serial port connection	<p>View the connected devices:</p> <pre>sastargets</pre>

The output shows the devices (disks and disk shelves) that the bridge is connected to. Output lines are sequentially numbered so that you can quickly count the devices. For example, the following output shows that 10 disks are connected:

```

Tgt VendorID ProductID      Type      SerialNumber
  0 NETAPP    X410_S15K6288A15 DISK      3QP1CLE300009940UHJV
  1 NETAPP    X410_S15K6288A15 DISK      3QP1ELF600009940V1BV
  2 NETAPP    X410_S15K6288A15 DISK      3QP1G3EW00009940U2M0
  3 NETAPP    X410_S15K6288A15 DISK      3QP1EWMP00009940U1X5
  4 NETAPP    X410_S15K6288A15 DISK      3QP1FZLE00009940G8YU
  5 NETAPP    X410_S15K6288A15 DISK      3QP1FZLF00009940TZKZ
  6 NETAPP    X410_S15K6288A15 DISK      3QP1CEB400009939MGXL
  7 NETAPP    X410_S15K6288A15 DISK      3QP1G7A900009939FNNT
  8 NETAPP    X410_S15K6288A15 DISK      3QP1FY0T00009940G8PA
  9 NETAPP    X410_S15K6288A15 DISK      3QP1FXW600009940VERQ

```



If the text “response truncated” appears at the beginning of the output, you can use Telnet to connect to the bridge and enter the same command to see all of the output.

- Verify that the command output shows that the bridge is connected to all disks and disk shelves in the stack that it is supposed to be connected to.

If the output is...	Then...
Correct	Repeat Step 1 for each remaining bridge.

Not correct	<p>a. Check for loose SAS cables or correct the SAS cabling by repeating the cabling.</p> <p>Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM12 modules</p> <p>b. Repeat Step 1.</p>
-------------	---

- Cable each bridge to the local FC switches, using the cabling in the table for your configuration and switch model and FC-to-SAS bridge model:



The second FC port connection on the FibreBridge 7500N bridge should not be cabled until zoning has been completed.

See the port assignments for your version of ONTAP.

- Repeat the previous step on the bridges at the partner site.

Related information

You need to verify that you are using the specified port assignments when you cable the FC switches.

[Port assignments for FC switches](#)

Secure or unsecure the FibreBridge bridge

To easily disable potentially unsecure Ethernet protocols on a bridge, beginning with ONTAP 9.5 you can secure the bridge. This disables the bridge's Ethernet ports. You can also reenable Ethernet access.

About this task

- Securing the bridge disables telnet and other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV) on the bridge.
- This procedure uses out-of-band management using the ONTAP prompt, which is available beginning with ONTAP 9.5.

You can issue the commands from the bridge CLI if you are not using out-of-band management.

- The `unsecurebridge` command can be used to re-enable the Ethernet ports.
- In ONTAP 9.7 and earlier, running the `securebridge` command on the ATTO FibreBridge might not update the bridge status correctly on the partner cluster. If this occurs, run the `securebridge` command from the partner cluster.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

- From the ONTAP prompt of the cluster containing the bridge, secure or unsecure the bridge.
 - The following command secures `bridge_A_1`:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command securebridge
```

- The following command unsecures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command unsecurebridge
```

2. From the ONTAP prompt of the cluster containing the bridge, save the bridge configuration:

```
storage bridge run-cli -bridge bridge-name -command saveconfiguration
```

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command  
saveconfiguration
```

3. From the ONTAP prompt of the cluster containing the bridge, restart the bridge's firmware:

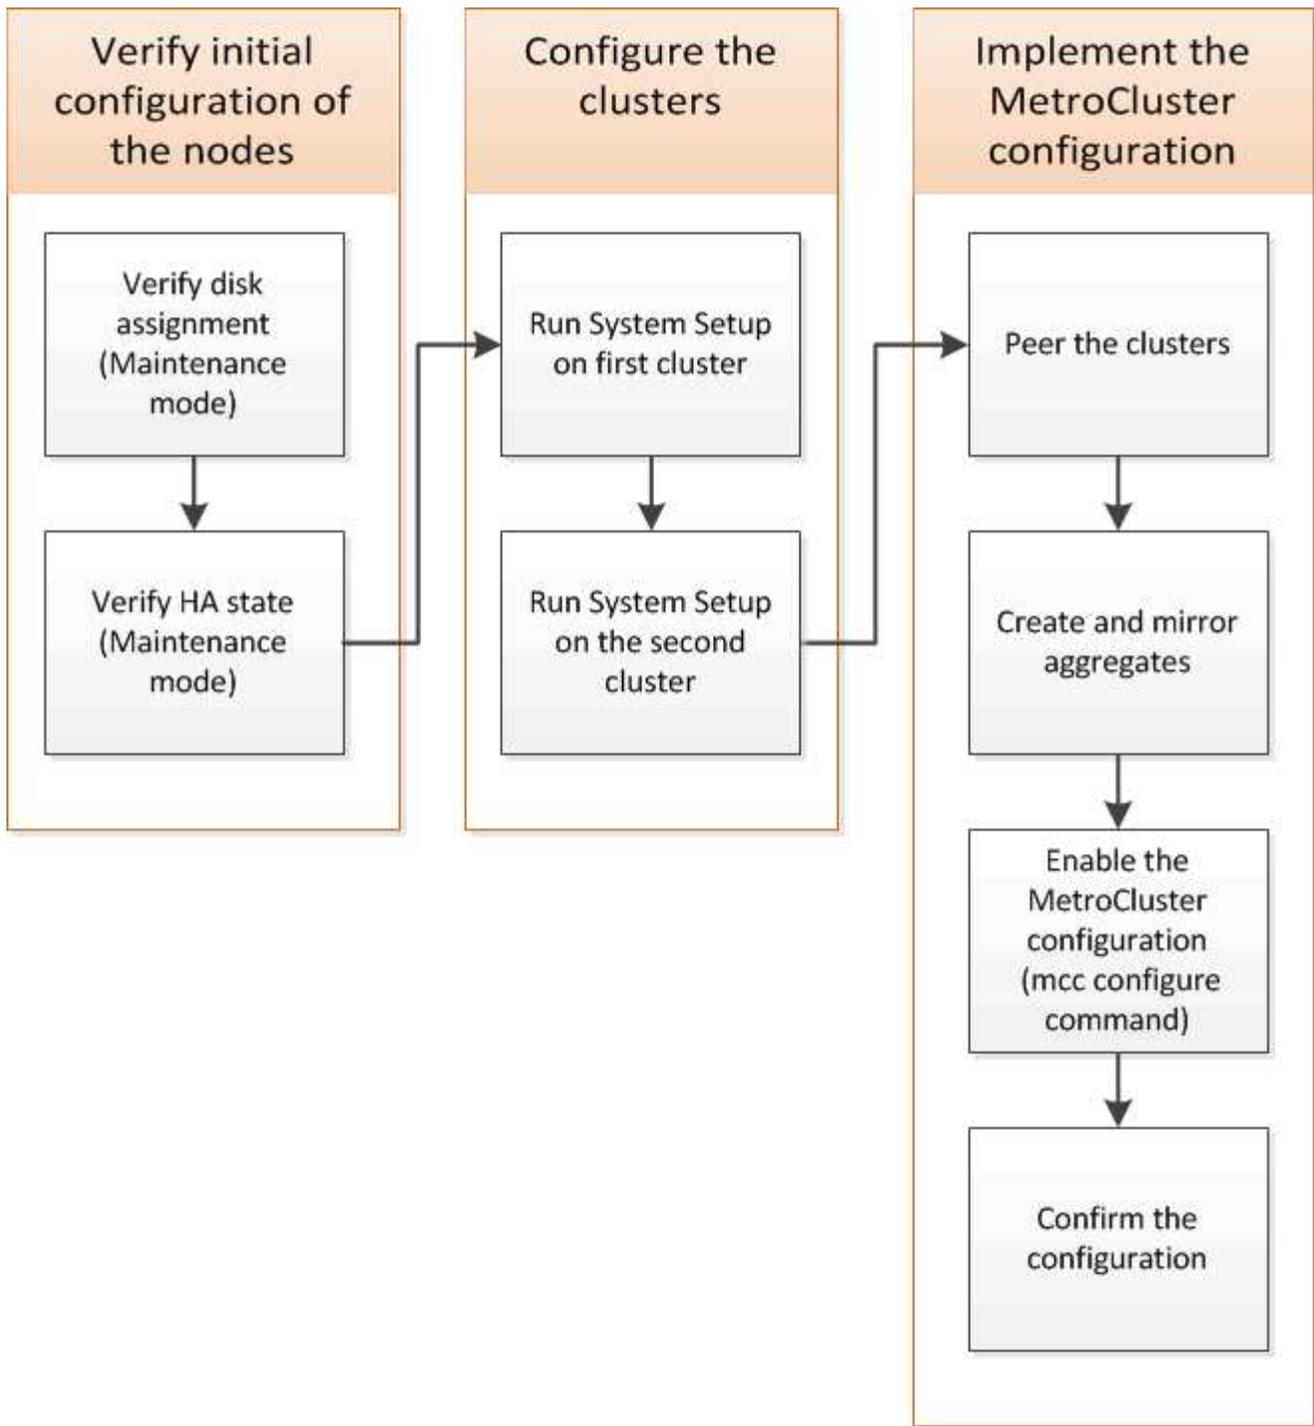
```
storage bridge run-cli -bridge bridge-name -command firmwarerestart
```

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command firmwarerestart
```

Configure the MetroCluster FC software in ONTAP

You must set up each node in the MetroCluster FC configuration in ONTAP, including the node-level configurations and the configuration of the nodes into two sites. You must also implement the MetroCluster relationship between the two sites.



Gathering required information

You need to gather the required IP addresses for the controller modules before you begin the configuration process.

IP network information worksheet for site A

You must obtain IP addresses and other network information for the first MetroCluster site (site A) from your network administrator before you configure the system.

Site A switch information (switched clusters)

When you cable the system, you need a host name and management IP address for each cluster switch. This information is not needed if you are using a two-node switchless cluster or have a two-node MetroCluster configuration (one node at each site).

Cluster switch	Host name	IP address	Network mask	Default gateway
Interconnect 1				
Interconnect 2				
Management 1				
Management 2				

Site A cluster creation information

When you first create the cluster, you need the following information:

Type of information	Your values
Cluster name Example used in this procedure: site_A	
DNS domain	
DNS name servers	
Location	
Administrator password	

Site A node information

For each node in the cluster, you need a management IP address, a network mask, and a default gateway.

Node	Port	IP address	Network mask	Default gateway
Node 1 Example used in this procedure: controller_A_1				

Node 2 Not required if using two-node MetroCluster configuration (one node at each site). Example used in this procedure: controller_A_2				
---	--	--	--	--

Site A LIFs and ports for cluster peering

For each node in the cluster, you need the IP addresses of two intercluster LIFs, including a network mask and a default gateway. The intercluster LIFs are used to peer the clusters.

Node	Port	IP address of intercluster LIF	Network mask	Default gateway
Node 1 IC LIF 1				
Node 1 IC LIF 2				
Node 2 IC LIF 1 Not required for two-node MetroCluster configurations (one node at each site).				
Node 2 IC LIF 2 Not required for two-node MetroCluster configurations (one node at each site).				

Site A time server information

You must synchronize the time, which requires one or more NTP time servers.

Node	Host name	IP address	Network mask	Default gateway
NTP server 1				
NTP server 2				

Site A AutoSupport information

You must configure AutoSupport on each node, which requires the following information:

Type of information		Your values
From email address		
Mail hosts	IP addresses or names	
Transport protocol	HTTP, HTTPS, or SMTP	
	Proxy server	
Recipient email addresses or distribution lists	Full-length messages	
	Concise messages	
	Partners	

Site A SP information

You must enable access to the Service Processor (SP) of each node for troubleshooting and maintenance, which requires the following network information for each node:

Node	IP address	Network mask	Default gateway
Node 1			
Node 2 Not required for two-node MetroCluster configurations (one node at each site).			

IP network information worksheet for Site B

You must obtain IP addresses and other network information for the second MetroCluster site (site B) from your network administrator before you configure the system.

Site B switch information (switched clusters)

When you cable the system, you need a host name and management IP address for each cluster switch. This information is not needed if you are using a two-node switchless cluster or you have a two-node MetroCluster configuration (one node at each site).

Cluster switch	Host name	IP address	Network mask	Default gateway
----------------	-----------	------------	--------------	-----------------

Interconnect 1				
Interconnect 2				
Management 1				
Management 2				

Site B cluster creation information

When you first create the cluster, you need the following information:

Type of information	Your values
Cluster name Example used in this procedure: site_B	
DNS domain	
DNS name servers	
Location	
Administrator password	

Site B node information

For each node in the cluster, you need a management IP address, a network mask, and a default gateway.

Node	Port	IP address	Network mask	Default gateway
Node 1 Example used in this procedure: controller_B_1				
Node 2 Not required for two-node MetroCluster configurations (one node at each site). Example used in this procedure: controller_B_2				

Site B LIFs and ports for cluster peering

For each node in the cluster, you need the IP addresses of two intercluster LIFs, including a network mask and a default gateway. The intercluster LIFs are used to peer the clusters.

Node	Port	IP address of intercluster LIF	Network mask	Default gateway
Node 1 IC LIF 1				
Node 1 IC LIF 2				
Node 2 IC LIF 1 Not required for two-node MetroCluster configurations (one node at each site).				
Node 2 IC LIF 2 Not required for two-node MetroCluster configurations (one node at each site).				

Site B time server information

You must synchronize the time, which requires one or more NTP time servers.

Node	Host name	IP address	Network mask	Default gateway
NTP server 1				
NTP server 2				

Site B AutoSupport information

You must configure AutoSupport on each node, which requires the following information:

Type of information		Your values
From email address		
Mail hosts	IP addresses or names	
Transport protocol	HTTP, HTTPS, or SMTP	
	Proxy server	

Recipient email addresses or distribution lists	Full-length messages	
	Concise messages	
	Partners	

Site B SP information

You must enable access to the Service Processor (SP) of each node for troubleshooting and maintenance, which requires the following network information for each node:

Node	IP address	Network mask	Default gateway
Node 1 (controller_B_1)			
Node 2 (controller_B_2)			
Not required for two-node MetroCluster configurations (one node at each site).			

Similarities and differences between standard cluster and MetroCluster configurations

The configuration of the nodes in each cluster in a MetroCluster configuration is similar to that of nodes in a standard cluster.

The MetroCluster configuration is built on two standard clusters. Physically, the configuration must be symmetrical, with each node having the same hardware configuration, and all of the MetroCluster components must be cabled and configured. However, the basic software configuration for nodes in a MetroCluster configuration is the same as that for nodes in a standard cluster.

Configuration step	Standard cluster configuration	MetroCluster configuration
Configure management, cluster, and data LIFs on each node.	Same in both types of clusters	
Configure the root aggregate.	Same in both types of clusters	
Configure nodes in the cluster as HA pairs	Same in both types of clusters	
Set up the cluster on one node in the cluster.	Same in both types of clusters	
Join the other node to the cluster.	Same in both types of clusters	
Create a mirrored root aggregate.	Optional	Required
Peer the clusters.	Optional	Required

Enable the MetroCluster configuration.	Does not apply	Required
--	----------------	----------

Verifying and configuring the HA state of components in Maintenance mode

When configuring a storage system in a MetroCluster FC configuration, you must make sure that the high-availability (HA) state of the controller module and chassis components is `mcc` or `mcc-2n` so that these components boot properly. Although this value should be preconfigured on systems received from the factory, you should still verify the setting before proceeding.

If the HA state of the controller module and chassis is incorrect, you cannot configure the MetroCluster without re-initializing the node. You must correct the setting using this procedure, and then initialize the system by using one of the following procedures:



- In a MetroCluster IP configuration, follow the steps in [Restore system defaults on a controller module](#).
- In a MetroCluster FC configuration, follow the steps in [Restore system defaults and configuring the HBA type on a controller module](#).

Before you begin

Verify that the system is in Maintenance mode.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The correct HA state depends on your MetroCluster configuration.

MetroCluster configuration type	HA state for all components...
Eight or four node MetroCluster FC configuration	<code>mcc</code>
Two-node MetroCluster FC configuration	<code>mcc-2n</code>
Eight or four node MetroCluster IP configuration	<code>mccip</code>

2. If the displayed system state of the controller is not correct, set the correct HA state for your configuration on the controller module:

MetroCluster configuration type	Command
Eight or four node MetroCluster FC configuration	<code>ha-config modify controller mcc</code>
Two-node MetroCluster FC configuration	<code>ha-config modify controller mcc-2n</code>
Eight or four node MetroCluster IP configuration	<code>ha-config modify controller mccip</code>

- If the displayed system state of the chassis is not correct, set the correct HA state for your configuration on the chassis:

MetroCluster configuration type	Command
Eight or four node MetroCluster FC configuration	<code>ha-config modify chassis mcc</code>
Two-node MetroCluster FC configuration	<code>ha-config modify chassis mcc-2n</code>
Eight or four node MetroCluster IP configuration	<code>ha-config modify chassis mccip</code>

- Boot the node to ONTAP:

```
boot_ontap
```

- Repeat this entire procedure to verify the HA state on each node in the MetroCluster configuration.

Restoring system defaults and configuring the HBA type on a controller module

About this task

To ensure a successful MetroCluster installation, reset and restore defaults on the controller modules.

Important

This task is only required for stretch configurations using FC-to-SAS bridges.

Steps

- At the LOADER prompt, return the environmental variables to their default setting:

```
set-defaults
```

- Boot the node into Maintenance mode, then configure the settings for any HBAs in the system:

- Boot into Maintenance mode:

```
boot_ontap maint
```

- Check the current settings of the ports:

```
ucadmin show
```

- Update the port settings as needed.

If you have this type of HBA and desired mode...	Use this command...
CNA FC	<code>ucadmin modify -m fc -t initiator <i>adapter_name</i></code>
CNA Ethernet	<code>ucadmin modify -mode cna <i>adapter_name</i></code>
FC target	<code>fcadmin config -t target <i>adapter_name</i></code>

FC initiator	<code>fcadmin config -t initiator adapter_name</code>
--------------	---

3. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

4. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

```
boot_ontap maint
```

5. Verify the changes you made:

If you have this type of HBA...	Use this command...
CNA	<code>ucadmin show</code>
FC	<code>fcadmin show</code>

6. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

7. Boot the node to the boot menu:

```
boot_ontap menu
```

After you run the command, wait until the boot menu is shown.

8. Clear the node configuration by typing “wipeconfig” at the boot menu prompt, and then press Enter.

The following screen shows the boot menu prompt:

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.

Selection (1-9)? wipeconfig

This option deletes critical system configuration, including cluster membership.

Warning: do not run this option on a HA node that has been taken over.

Are you sure you want to continue?: yes

Rebooting to finish wipeconfig request.

Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems

If you are using the X1132A-R6 quad-port card on a FAS8020 system, you can enter Maintenance mode to configure the 1a and 1b ports for FC-VI and initiator usage. This is not required on MetroCluster systems received from the factory, in which the ports are set appropriately for your configuration.

About this task

This task must be performed in Maintenance mode.



Converting an FC port to an FC-VI port with the `ucadmin` command is only supported on the FAS8020 and AFF 8020 systems. Converting FC ports to FCVI ports is not supported on any other platform.

Steps

1. Disable the ports:

```
storage disable adapter 1a
```

```
storage disable adapter 1b
```

```

*> storage disable adapter 1a
Jun 03 02:17:57 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1a.
Host adapter 1a disable succeeded
Jun 03 02:17:57 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1a is now offline.
*> storage disable adapter 1b
Jun 03 02:18:43 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1b.
Host adapter 1b disable succeeded
Jun 03 02:18:43 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1b is now offline.
*>

```

2. Verify that the ports are disabled:

```
ucadmin show
```

```

*> ucadmin show

```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
...					
1a	fc	initiator	-	-	offline
1b	fc	initiator	-	-	offline
1c	fc	initiator	-	-	online
1d	fc	initiator	-	-	online

3. Set the a and b ports to FC-VI mode:

```
ucadmin modify -adapter 1a -type fcvi
```

The command sets the mode on both ports in the port pair, 1a and 1b (even though only 1a is specified in the command).

```

*> ucadmin modify -t fcvi 1a
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1a. Reboot the controller for the changes to
take effect.
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1b. Reboot the controller for the changes to
take effect.

```

4. Confirm that the change is pending:

```
ucadmin show
```

```
*> ucadmin show
      Current   Current   Pending   Pending   Admin
Adapter Mode     Type     Mode     Type     Status
-----
...
1a    fc      initiator -        fcvi     offline
1b    fc      initiator -        fcvi     offline
1c    fc      initiator -         -        online
1d    fc      initiator -         -        online
```

- 5. Shut down the controller, and then reboot into Maintenance mode.
- 6. Confirm the configuration change:

```
ucadmin show local
```

```
Node           Adapter  Mode     Type     Mode     Type     Status
-----
...
controller_B_1 1a       fc       fcvi     -        -        online
controller_B_1 1b       fc       fcvi     -        -        online
controller_B_1 1c       fc       initiator -        -        online
controller_B_1 1d       fc       initiator -        -        online
6 entries were displayed.
```

Verifying disk assignment in Maintenance mode in an eight-node or a four-node configuration

Before fully booting the system to ONTAP, you can optionally boot to Maintenance mode and verify the disk assignment on the nodes. The disks should be assigned to create a fully symmetric active-active configuration, where each pool has an equal number of disks assigned to them.

About this task

New MetroCluster systems have disk assignment completed prior to shipment.

The following table shows example pool assignments for a MetroCluster configuration. Disks are assigned to pools on a per-shelf basis.

Disk shelves at Site A

Disk shelf (sample_shelf_name)...	Belongs to...	And is assigned to that node's...
Disk shelf 1 (shelf_A_1_1)	Node A 1	Pool 0
Disk shelf 2 (shelf_A_1_3)		
Disk shelf 3 (shelf_B_1_1)	Node B 1	Pool 1
Disk shelf 4 (shelf_B_1_3)		
Disk shelf 5 (shelf_A_2_1)	Node A 2	Pool 0
Disk shelf 6 (shelf_A_2_3)		
Disk shelf 7 (shelf_B_2_1)	Node B 2	Pool 1
Disk shelf 8 (shelf_B_2_3)		
Disk shelf 1 (shelf_A_3_1)	Node A 3	Pool 0
Disk shelf 2 (shelf_A_3_3)		
Disk shelf 3 (shelf_B_3_1)	Node B 3	Pool 1
Disk shelf 4 (shelf_B_3_3)		
Disk shelf 5 (shelf_A_4_1)	Node A 4	Pool 0
Disk shelf 6 (shelf_A_4_3)		
Disk shelf 7 (shelf_B_4_1)	Node B 4	Pool 1
Disk shelf 8 (shelf_B_4_3)		

Disk shelves at Site B

Disk shelf (sample_shelf_name)...	Belongs to...	And is assigned to that node's...
Disk shelf 9 (shelf_B_1_2)	Node B 1	Pool 0
Disk shelf 10 (shelf_B_1_4)		
Disk shelf 11 (shelf_A_1_2)	Node A 1	Pool 1
Disk shelf 12 (shelf_A_1_4)		
Disk shelf 13 (shelf_B_2_2)	Node B 2	Pool 0
Disk shelf 14 (shelf_B_2_4)		
Disk shelf 15 (shelf_A_2_2)	Node A 2	Pool 1
Disk shelf 16 (shelf_A_2_4)		

Disk shelf 1 (shelf_B_3_2)	Node A 3	Pool 0
Disk shelf 2 (shelf_B_3_4)		
Disk shelf 3 (shelf_A_3_2)	Node B 3	Pool 1
Disk shelf 4 (shelf_A_3_4)		
Disk shelf 5 (shelf_B_4_2)	Node A 4	Pool 0
Disk shelf 6 (shelf_B_4_4)		
Disk shelf 7 (shelf_A_4_2)	Node B 4	Pool 1
Disk shelf 8 (shelf_A_4_4)		

Steps

1. Confirm the shelf assignments:

```
disk show -v
```

2. If necessary, explicitly assign disks on the attached disk shelves to the appropriate pool:

```
disk assign
```

Using wildcards in the command enables you to assign all of the disks on a disk shelf with one command. You can identify the disk shelf IDs and bays for each disk with the `storage show disk -x` command.

Assigning disk ownership in non-AFF systems

If the MetroCluster nodes do not have the disks correctly assigned, or if you are using DS460C disk shelves in your configuration, you must assign disks to each of the nodes in the MetroCluster configuration on a shelf-by-shelf basis. You will create a configuration in which each node has the same number of disks in its local and remote disk pools.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

If your configuration does not include DS460C disk shelves, this task is not required if disks were correctly assigned when received from the factory.



Pool 0 always contains the disks that are found at the same site as the storage system that owns them.

Pool 1 always contains the disks that are remote to the storage system that owns them.

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

Assign these disks in the drawer...	To this node and pool...
0 - 2	Local node's pool 0
3 - 5	HA partner node's pool 0
6 - 8	DR partner of the local node's pool 1
9 - 11	DR partner of the HA partner's pool 1

This disk assignment pattern ensures that an aggregate is minimally affected in case a drawer goes offline.

Steps

1. If you have not done so, boot each system into Maintenance mode.
2. Assign the disk shelves to the nodes located at the first site (site A):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

If storage controller Controller_A_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0

*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1
*> disk assign -shelf FC_switch_B_1:1-4.shelf2 -p 1
```

- b. Repeat the process for the second node at the local site, systematically assigning the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

If storage controller Controller_A_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1

*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1
```

3. Assign the disk shelves to the nodes located at the second site (site B):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node at the remote site, systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf local-switch-nameshelf-name -p pool
```

If storage controller Controller_B_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1
```

- b. Repeat the process for the second node at the remote site, systematically assigning its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf shelf-name -p pool
```

If storage controller Controller_B_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1
```

4. Confirm the shelf assignments:

```
storage show shelf
```

5. Exit Maintenance mode:

```
halt
```

6. Display the boot menu:

```
boot_ontap menu
```

7. On each node, select option 4 to initialize all disks.

Assigning disk ownership in AFF systems

If you are using AFF systems in a configuration with mirrored aggregates and the nodes do not have the disks (SSDs) correctly assigned, you should assign half the disks on each shelf to one local node and the other half of the disks to its HA partner node. You should create a configuration in which each node has the same number of disks in its local and remote disk pools.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

This does not apply to configurations which have unmirrored aggregates, an active/passive configuration, or that have an unequal number of disks in local and remote pools.

This task is not required if disks were correctly assigned when received from the factory.



Pool 0 always contains the disks that are found at the same site as the storage system that owns them.

Pool 1 always contains the disks that are remote to the storage system that owns them.

Steps

1. If you have not done so, boot each system into Maintenance mode.
2. Assign the disks to the nodes located at the first site (site A):

You should assign an equal number of disks to each pool.

- a. On the first node, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0:

```
disk assign -shelf <shelf-name> -p <pool> -n <number-of-disks>
```

If storage controller Controller_A_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0 -n 4

*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf2 -p 1 -n 4
```

- b. Repeat the process for the second node at the local site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1:

```
disk assign -disk disk-name -p pool
```

If storage controller Controller_A_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1 -n 4

*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1 -n 4
```

3. Assign the disks to the nodes located at the second site (site B):

You should assign an equal number of disks to each pool.

- a. On the first node at the remote site, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0:

```
disk assign -disk disk-name -p pool
```

If storage controller Controller_B_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0 -n 4

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1 -n 4
```

- b. Repeat the process for the second node at the remote site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1:

```
disk assign -disk disk-name -p pool
```

If storage controller Controller_B_2 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0 -n 4

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1 -n 4
```

4. Confirm the disk assignments:

```
storage show disk
```

5. Exit Maintenance mode:

```
halt
```

6. Display the boot menu:

boot_ontap menu

7. On each node, select option **4** to initialize all disks.

Verifying disk assignment in Maintenance mode in a two-node configuration

Before fully booting the system to ONTAP, you can optionally boot the system to Maintenance mode and verify the disk assignment on the nodes. The disks should be assigned to create a fully symmetric configuration with both sites owning their own disk shelves and serving data, where each node and each pool have an equal number of mirrored disks assigned to them.

Before you begin

The system must be in Maintenance mode.

About this task

New MetroCluster systems have disk assignment completed prior to shipment.

The following table shows example pool assignments for a MetroCluster configuration. Disks are assigned to pools on a per-shelf basis.

Disk shelf (example name)...	At site...	Belongs to...	And is assigned to that node's...
Disk shelf 1 (shelf_A_1_1)	Site A	Node A 1	Pool 0
Disk shelf 2 (shelf_A_1_3)			
Disk shelf 3 (shelf_B_1_1)		Node B 1	Pool 1
Disk shelf 4 (shelf_B_1_3)			
Disk shelf 9 (shelf_B_1_2)	Site B	Node B 1	Pool 0
Disk shelf 10 (shelf_B_1_4)			
Disk shelf 11 (shelf_A_1_2)		Node A 1	Pool 1
Disk shelf 12 (shelf_A_1_4)			

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

Assign these disks in the drawer...	To this node and pool...
1 - 6	Local node's pool 0
7 - 12	DR partner's pool 1

This disk assignment pattern minimizes the effect on an aggregate if a drawer goes offline.

Steps

1. If your system was received from the factory, confirm the shelf assignments:

```
disk show -v
```

2. If necessary, you can explicitly assign disks on the attached disk shelves to the appropriate pool by using the `disk assign` command.

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1. You should assign an equal number of shelves to each pool.

- a. If you have not done so, boot each system into Maintenance mode.
- b. On the node on site A, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf disk_shelf_name -p pool
```

If storage controller node_A_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf shelf_A_1_1 -p 0
*> disk assign -shelf shelf_A_1_3 -p 0

*> disk assign -shelf shelf_A_1_2 -p 1
*> disk assign -shelf shelf_A_1_4 -p 1
```

- c. On the node at the remote site (site B), systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf disk_shelf_name -p pool
```

If storage controller node_B_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf shelf_B_1_2 -p 0
*> disk assign -shelf shelf_B_1_4 -p 0

*> disk assign -shelf shelf_B_1_1 -p 1
*> disk assign -shelf shelf_B_1_3 -p 1
```

- d. Show the disk shelf IDs and bays for each disk:

```
disk show -v
```

Setting up ONTAP

You must set up ONTAP on each controller module.

If you need to netboot the new controllers, see [Netbooting the new controller modules](#) in the *MetroCluster Upgrade, Transition, and Expansion Guide*.

Choices

- [Setting up ONTAP in a two-node MetroCluster configuration](#)

- [Setting up ONTAP in an eight-mode or four-node MetroCluster configuration](#)

Setting up ONTAP in a two-node MetroCluster configuration

In a two-node MetroCluster configuration, on each cluster you must boot up the node, exit the Cluster Setup wizard, and use the cluster setup command to configure the node into a single-node cluster.

Before you begin

You must not have configured the Service Processor.

About this task

This task is for two-node MetroCluster configurations using native NetApp storage.

This task must be performed on both clusters in the MetroCluster configuration.

For more general information about setting up ONTAP, see [Set up ONTAP](#).

Steps

1. Power on the first node.



You must repeat this step on the node at the disaster recovery (DR) site.

The node boots, and then the Cluster Setup wizard starts on the console, informing you that AutoSupport will be enabled automatically.

```
::> Welcome to the cluster setup wizard.
```

You can enter the following commands at any time:

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.
```

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
resolution, should a problem occur on your system.
For further information on AutoSupport, see:
<http://support.netapp.com/autosupport/>

```
Type yes to confirm and continue {yes}: yes
```

```
Enter the node management interface port [e0M]:
```

```
Enter the node management interface IP address [10.101.01.01]:
```

```
Enter the node management interface netmask [101.010.101.0]:
```

```
Enter the node management interface default gateway [10.101.01.0]:
```

```
Do you want to create a new cluster or join an existing cluster?  
{create, join}:
```

2. Create a new cluster:

```
create
```

3. Choose whether the node is to be used as a single node cluster.

```
Do you intend for this node to be used as a single node cluster? {yes,  
no} [yes]:
```

4. Accept the system default `yes` by pressing Enter, or enter your own values by typing `no`, and then pressing

Enter.

5. Follow the prompts to complete the **Cluster Setup** wizard, pressing Enter to accept the default values or typing your own values and then pressing Enter.

The default values are determined automatically based on your platform and network configuration.

6. After you complete the **Cluster Setup** wizard and it exits, verify that the cluster is active and the first node is healthy: `

```
cluster show
```

The following example shows a cluster in which the first node (cluster1-01) is healthy and eligible to participate:

```
cluster1::> cluster show
Node                Health  Eligibility
-----
cluster1-01        true   true
```

If it becomes necessary to change any of the settings you entered for the admin SVM or node SVM, you can access the Cluster Setup wizard by using the cluster setup command.

Setting up ONTAP in an eight-node or four-node MetroCluster configuration

After you boot each node, you are prompted to run the System Setup tool to perform basic node and cluster configuration. After configuring the cluster, you return to the ONTAP CLI to create aggregates and create the MetroCluster configuration.

Before you begin

You must have cabled the MetroCluster configuration.

About this task

This task is for eight-node or four-node MetroCluster configurations using native NetApp storage.

New MetroCluster systems are preconfigured; you do not need to perform these steps. However, you should configure the AutoSupport tool.

This task must be performed on both clusters in the MetroCluster configuration.

This procedure uses the System Setup tool. If desired, you can use the CLI cluster setup wizard instead.

Steps

1. If you have not already done so, power up each node and let them boot completely.

If the system is in Maintenance mode, issue the halt command to exit Maintenance mode, and then issue the following command from the LOADER prompt:

```
boot_ontap
```

The output should be similar to the following:

```
Welcome to node setup
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the setup wizard.
```

```
Any changes you made before quitting will be saved.
```

```
To accept a default or omit a question, do not enter a value.
```

```
.  
. .  
. .
```

2. Enable the AutoSupport tool by following the directions provided by the system.
3. Respond to the prompts to configure the node management interface.

The prompts are similar to the following:

```
Enter the node management interface port: [e0M]:  
Enter the node management interface IP address: 10.228.160.229  
Enter the node management interface netmask: 225.225.252.0  
Enter the node management interface default gateway: 10.228.160.1
```

4. Confirm that nodes are configured in high-availability mode:

```
storage failover show -fields mode
```

If not, you must issue the following command on each node and reboot the node:

```
storage failover modify -mode ha -node localhost
```

This command configures high availability mode but does not enable storage failover. Storage failover is automatically enabled when the MetroCluster configuration is performed later in the configuration process.

5. Confirm that you have four ports configured as cluster interconnects:

```
network port show
```

The following example shows output for cluster_A:

```

cluster_A::> network port show

```

(Mbps)						Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper

node_A_1						
	**e0a	Cluster	Cluster	up	1500	
	auto/1000					
	e0b	Cluster	Cluster	up	1500	
	auto/1000**					
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000
node_A_2						
	**e0a	Cluster	Cluster	up	1500	
	auto/1000					
	e0b	Cluster	Cluster	up	1500	
	auto/1000**					
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000

14 entries were displayed.

6. If you are creating a two-node switchless cluster (a cluster without cluster interconnect switches), enable the switchless-cluster networking mode:

a. Change to the advanced privilege level:

```
set -privilege advanced
```

You can respond `y` when prompted to continue into advanced mode. The advanced mode prompt appears (`*>`).

b. Enable switchless-cluster mode:

```
network options switchless-cluster modify -enabled true
```

c. Return to the admin privilege level:

```
set -privilege admin
```

7. Launch the System Setup tool as directed by the information that appears on the system console after the initial boot.

8. Use the System Setup tool to configure each node and create the cluster, but do not create aggregates.



You create mirrored aggregates in later tasks.

After you finish

Return to the ONTAP command-line interface and complete the MetroCluster configuration by performing the tasks that follow.

Configuring the clusters into a MetroCluster configuration

You must peer the clusters, mirror the root aggregates, create a mirrored data aggregate, and then issue the command to implement the MetroCluster operations.

About this task

Before you run `metrocluster configure`, HA mode and DR mirroring are not enabled and you might see an error message related to this expected behavior. You enable HA mode and DR mirroring later when you run the command `metrocluster configure` to implement the configuration.

Peering the clusters

The clusters in the MetroCluster configuration must be in a peer relationship so that they can communicate with each other and perform the data mirroring essential to MetroCluster disaster recovery.

Configuring intercluster LIFs

You must create intercluster LIFs on ports used for communication between the MetroCluster partner clusters. You can use dedicated ports or ports that also have data traffic.

Choices

- [Configuring intercluster LIFs on dedicated ports](#)
- [Configuring intercluster LIFs on shared data ports](#)

Configuring intercluster LIFs on dedicated ports

You can configure intercluster LIFs on dedicated ports. Doing so typically increases the available bandwidth for replication traffic.

Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in "cluster01":

```
cluster01::> network port show
```

(Mbps)							Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	

cluster01-01							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	
	e0e	Default	Default	up	1500	auto/1000	
	e0f	Default	Default	up	1500	auto/1000	
cluster01-02							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	
	e0e	Default	Default	up	1500	auto/1000	
	e0f	Default	Default	up	1500	auto/1000	

2. Determine which ports are available to dedicate to intercluster communication:

```
network interface show -fields home-port,curr-port
```

For complete command syntax, see the man page.

The following example shows that ports "e0e" and "e0f" have not been assigned LIFs:

```
cluster01::> network interface show -fields home-port,curr-port
```

vserver	lif	home-port	curr-port
Cluster	cluster01-01_clus1	e0a	e0a
Cluster	cluster01-01_clus2	e0b	e0b
Cluster	cluster01-02_clus1	e0a	e0a
Cluster	cluster01-02_clus2	e0b	e0b
cluster01	cluster_mgmt	e0c	e0c
cluster01	cluster01-01_mgmt1	e0c	e0c
cluster01	cluster01-02_mgmt1	e0c	e0c

3. Create a failover group for the dedicated ports:

```
network interface failover-groups create -vserver system_SVM -failover-group
failover_group -targets physical_or_logical_ports
```

The following example assigns ports "e0e" and "e0f" to the failover group intercluster01 on the system "SVMcluster01":

```
cluster01::> network interface failover-groups create -vserver cluster01
-failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f
```

4. Verify that the failover group was created:

```
network interface failover-groups show
```

For complete command syntax, see the man page.

```
cluster01::> network interface failover-groups show
```

Vserver	Group	Failover Targets
-----	-----	-----
Cluster	Cluster	cluster01-01:e0a, cluster01-01:e0b, cluster01-02:e0a, cluster01-02:e0b
cluster01	Default	cluster01-01:e0c, cluster01-01:e0d, cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f
	intercluster01	cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

ONTAP 9.6 and later

```
network interface create -vserver system_SVM -lif LIF_name -service-policy
default-intercluster -home-node node -home-port port -address port_IP
-netmask netmask -failover-group failover_group
```

ONTAP 9.5 and earlier

```
network interface create -vserver system_SVM -lif LIF_name -role
intercluster -home-node node -home-port port -address port_IP -netmask
netmask -failover-group failover_group
```

For complete command syntax, see the man page.

The following example creates intercluster LIFs "cluster01_icl01" and "cluster01_icl02" in the failover group "intercluster01":

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0 -failover-group intercluster01
```

6. Verify that the intercluster LIFs were created:

ONTAP 9.6 and later

Run the command: `network interface show -service-policy default-intercluster`

ONTAP 9.5 and earlier

Run the command: `network interface show -role intercluster`

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-intercluster
          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
cluster01
          cluster01_icl01
                up/up      192.168.1.201/24  cluster01-01  e0e
true
          cluster01_icl02
                up/up      192.168.1.202/24  cluster01-02  e0f
true

```

7. Verify that the intercluster LIFs are redundant:

ONTAP 9.6 and later

Run the command: `network interface show -service-policy default-intercluster -failover`

ONTAP 9.5 and earlier

Run the command: `network interface show -role intercluster -failover`

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01_icl01" and "cluster01_icl02" on the SVM "e0e" port will fail over to the "e0f" port.

```

cluster01::> network interface show -service-policy default-intercluster
-failover
          Logical      Home      Failover      Failover
Vserver   Interface  Node:Port  Policy      Group
-----
-----
cluster01
          cluster01_icl01  cluster01-01:e0e  local-only
intercluster01
                Failover Targets:  cluster01-01:e0e,
                cluster01-01:e0f
          cluster01_icl02  cluster01-02:e0e  local-only
intercluster01
                Failover Targets:  cluster01-02:e0e,
                cluster01-02:e0f

```

Related information

[Considerations when using dedicated ports](#)

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Configuring intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```
cluster01::> network port show
```

							Speed
(Mbps)							
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	
-----	-----	-----	-----	-----	-----	-----	
cluster01-01							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	
cluster01-02							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	

2. Create intercluster LIFs on the system SVM:

ONTAP 9.6 and later

Run the command: `network interface create -vserver system_SVM -lif LIF_name -service-policy default-intercluster -home-node node -home-port port -address port_IP -netmask netmask`

ONTAP 9.5 and earlier

Run the command:

```
network interface create -vserver system_SVM -lif LIF_name -role
intercluster -home-node node -home-port port -address port_IP -netmask
netmask
```

For complete command syntax, see the man page.

The following example creates intercluster LIFs `cluster01_icl01` and `cluster01_icl02`:

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0c
-address 192.168.1.201
-netmask 255.255.255.0
```

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0c
-address 192.168.1.202
-netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

ONTAP 9.6 and later

Run the command: `network interface show -service-policy default-intercluster`

ONTAP 9.5 and earlier

Run the command: `network interface show -role intercluster`

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-intercluster
          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node        Port
Home
-----
-----
cluster01
          cluster01_icl01
                up/up      192.168.1.201/24  cluster01-01  e0c
true
          cluster01_icl02
                up/up      192.168.1.202/24  cluster01-02  e0c
true

```

4. Verify that the intercluster LIFs are redundant:

ONTAP 9.6 and later

Run the command: `network interface show -service-policy default-intercluster -failover`

ONTAP 9.5 and earlier

Run the command:

`network interface show -role intercluster -failover`

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01_icl01" and "cluster01_icl02" on the "e0c" port will fail over to the "e0d" port.

```

cluster01::> network interface show -service-policy default-intercluster
-failover
          Logical      Home      Failover      Failover
Vserver   Interface  Node:Port  Policy        Group
-----
-----
cluster01
          cluster01_icl01  cluster01-01:e0c  local-only
192.168.1.201/24
                                Failover Targets: cluster01-01:e0c,
                                                cluster01-01:e0d
          cluster01_icl02  cluster01-02:e0c  local-only
192.168.1.201/24
                                Failover Targets: cluster01-02:e0c,
                                                cluster01-02:e0d

```

Related information

[Considerations when sharing data ports](#)

Creating a cluster peer relationship

You must create the cluster peer relationship between the MetroCluster clusters.

About this task

You can use the `cluster peer create` command to create a peer relationship between a local and remote cluster. After the peer relationship has been created, you can run `cluster peer create` on the remote cluster to authenticate it to the local cluster.

Before you begin

- You must have created intercluster LIFs on every node in the clusters that are being peered.
- The clusters must be running ONTAP 9.3 or later.

Steps

1. On the destination cluster, create a peer relationship with the source cluster:

```
cluster peer create -generate-passphrase -offer-expiration MM/DD/YYYY  
HH:MM:SS|1...7days|1...168hours -peer-addr peer_LIF_IPs -ip-space ip-space
```

If you specify both `-generate-passphrase` and `-peer-addr`, only the cluster whose intercluster LIFs are specified in `-peer-addr` can use the generated password.

You can ignore the `-ip-space` option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship on an unspecified remote cluster:

```
cluster02::> cluster peer create -generate-passphrase -offer-expiration  
2days  
  
                Passphrase: UCa+6lRVICXeL/gq1WrK7ShR  
                Expiration Time: 6/7/2017 08:16:10 EST  
Initial Allowed Vserver Peers: -  
                Intercluster LIF IP: 192.140.112.101  
                Peer Cluster Name: Clus_7ShR (temporary generated)  
  
Warning: make a note of the passphrase - it cannot be displayed again.
```

2. On the source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addr peer_LIF_IPs -ip-space ip-space
```

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses "192.140.112.101" and "192.140.112.102":

```
cluster01::> cluster peer create -peer-addr  
192.140.112.101,192.140.112.102
```

Notice: Use a generated passphrase or choose a passphrase of 8 or more characters.

To ensure the authenticity of the peering relationship, use a phrase or sequence of characters that would be hard to guess.

```
Enter the passphrase:  
Confirm the passphrase:
```

```
Clusters cluster02 and cluster01 are peered.
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

```
cluster01::> cluster peer show -instance  
  
Peer Cluster Name: cluster02  
Remote Intercluster Addresses: 192.140.112.101,  
192.140.112.102  
Availability of the Remote Cluster: Available  
Remote Cluster Name: cluster2  
Active IP Addresses: 192.140.112.101,  
192.140.112.102  
  
Cluster Serial Number: 1-80-123456  
Address Family of Relationship: ipv4  
Authentication Status Administrative: no-authentication  
Authentication Status Operational: absent  
Last Update Time: 02/05 21:05:41  
IPspace for the Relationship: Default
```

4. Check the connectivity and status of the nodes in the peer relationship:

```
cluster peer health show
```

```

cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
-----
cluster01-01
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
cluster01-02
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true

```

Creating a cluster peer relationship (ONTAP 9.2 and earlier)

You can use the `cluster peer create` command to initiate a request for a peering relationship between a local and remote cluster. After the peer relationship has been requested by the local cluster, you can run `cluster peer create` on the remote cluster to accept the relationship.

Before you begin

- You must have created intercluster LIFs on every node in the clusters being peered.
- The cluster administrators must have agreed on the passphrase that each cluster will use to authenticate itself to the other.

Steps

1. On the data protection destination cluster, create a peer relationship with the data protection source cluster:

```
cluster peer create -peer-addr peer_LIF_IPs -ipspace ipspace
```

You can ignore the `-ipspace` option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship with the remote cluster at intercluster LIF IP addresses "192.168.2.201" and "192.168.2.202":

```

cluster02::> cluster peer create -peer-addr 192.168.2.201,192.168.2.202
Enter the passphrase:
Please enter the passphrase again:

```

Enter the passphrase for the peer relationship when prompted.

2. On the data protection source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addr peer_LIF_IPs -ip-space ip-space
```

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses "192.140.112.203" and "192.140.112.204":

```
cluster01::> cluster peer create -peer-addr 192.168.2.203,192.168.2.204
Please confirm the passphrase:
Please confirm the passphrase again:
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

For complete command syntax, see the man page.

```
cluster01::> cluster peer show -instance
Peer Cluster Name: cluster01
Remote Intercluster Addresses: 192.168.2.201,192.168.2.202
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.168.2.201,192.168.2.202
Cluster Serial Number: 1-80-000013
```

4. Check the connectivity and status of the nodes in the peer relationship:

```
cluster peer health show`
```

For complete command syntax, see the man page.

```

cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
-----
cluster01-01
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
cluster01-02
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true

```

Mirroring the root aggregates

You must mirror the root aggregates to provide data protection.

About this task

By default, the root aggregate is created as RAID-DP type aggregate. You can change the root aggregate from RAID-DP to RAID4 type aggregate. The following command modifies the root aggregate for RAID4 type aggregate:

```
storage aggregate modify -aggregate aggr_name -raidtype raid4
```



On non-ADP systems, the RAID type of the aggregate can be modified from the default RAID-DP to RAID4 before or after the aggregate is mirrored.

Steps

1. Mirror the root aggregate:

```
storage aggregate mirror aggr_name
```

The following command mirrors the root aggregate for controller_A_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote

MetroCluster site.

2. Repeat the previous step for each node in the MetroCluster configuration.

Related information

[Logical storage management with the CLI](#)

Creating a mirrored data aggregate on each node

You must create a mirrored data aggregate on each node in the DR group.

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.
- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
See [Disk and aggregate management](#).
- Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.

Steps

1. Display a list of available spares:

```
storage disk show -spare -owner node_name
```

2. Create the aggregate by using the `storage aggregate create -mirror true` command.

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To ensure that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include



In the minimum-supported configuration, in which a limited number of drives are available, you must use the `force-small-aggregate` option to allow the creation of a three disk RAID-DP aggregate.

- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group

- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create` man page.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10 -node
node_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_node_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate aggregate-name
```

Creating unmirrored data aggregates

You can optionally create unmirrored data aggregates for data that does not require the redundant mirroring provided by MetroCluster configurations.

Before you begin

- Verify that you know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can verify that the correct drive type is selected.



In MetroCluster FC configurations, the unmirrored aggregates will only be online after a switchover if the remote disks in the aggregate are accessible. If the ISLs fail, the local node may be unable to access the data in the unmirrored remote disks. The failure of an aggregate can lead to a reboot of the local node.

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.



The unmirrored aggregates must be local to the node owning them.

- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- *Disks and aggregates management* contains more information about mirroring aggregates.

Steps

1. Display a list of available spares:

```
storage disk show -spare -owner node_name
```

2. Create the aggregate:

```
storage aggregate create
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To verify that the aggregate is created on a specific node, you should use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include
- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create man` page.

The following command creates a unmirrored aggregate with 10 disks:

```
controller_A_1::> storage aggregate create aggr1_controller_A_1
-diskcount 10 -node controller_A_1
[Job 15] Job is queued: Create aggr1_controller_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate aggregate-name
```

Related information

[Disk and tier \(aggregate\) management](#)

Implementing the MetroCluster configuration

You must run the `metrocluster configure` command to start data protection in a MetroCluster configuration.

Before you begin

- There should be at least two non-root mirrored data aggregates on each cluster.

Additional data aggregates can be either mirrored or unmirrored.

You can verify this with the `storage aggregate show` command.



If you want to use a single mirrored data aggregate, then see [Step 1](#) for instructions.

- The ha-config state of the controllers and chassis must be "mcc".

About this task

You issue the `metrocluster configure` command once, on any of the nodes, to enable the MetroCluster configuration. You do not need to issue the command on each of the sites or nodes, and it does not matter which node or site you choose to issue the command on.

The `metrocluster configure` command automatically pairs the two nodes with the lowest system IDs in each of the two clusters as disaster recovery (DR) partners. In a four-node MetroCluster configuration, there are two DR partner pairs. The second DR pair is created from the two nodes with higher system IDs.



You must **not** configure Onboard Key Manager (OKM) or external key management before you run the command `metrocluster configure`.

Steps

1. Configure the MetroCluster in the following format:

If your MetroCluster configuration has...	Then do this...
Multiple data aggregates	From any node's prompt, configure MetroCluster: <code>metrocluster configure node-name</code>
A single mirrored data aggregate	<p>a. From any node's prompt, change to the advanced privilege level:</p> <pre>set -privilege advanced</pre> <p>You need to respond with <code>y</code> when you are prompted to continue into advanced mode and you see the advanced mode prompt (<code>*></code>).</p> <p>b. Configure the MetroCluster with the <code>-allow -with-one-aggregate true</code> parameter:</p> <pre>metrocluster configure -allow-with-one-aggregate true node-name</pre> <p>c. Return to the admin privilege level:</p> <pre>set -privilege admin</pre>



The best practice is to have multiple data aggregates. If the first DR group has only one aggregate, and you want to add a DR group with one aggregate, you must move the metadata volume off the single data aggregate. For more information on this procedure, see [Moving a metadata volume in MetroCluster configurations](#).

The following command enables the MetroCluster configuration on all of the nodes in the DR group that

contains controller_A_1:

```
cluster_A::*> metrocluster configure -node-name controller_A_1  
  
[Job 121] Job succeeded: Configure is successful.
```

2. Verify the networking status on site A:

```
network port show
```

The following example shows the network port usage on a four-node MetroCluster configuration:

```
cluster_A::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper
controller_A_1						
	e0a	Cluster	Cluster	up	9000	auto/1000
	e0b	Cluster	Cluster	up	9000	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000
controller_A_2						
	e0a	Cluster	Cluster	up	9000	auto/1000
	e0b	Cluster	Cluster	up	9000	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000

14 entries were displayed.

3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Verify the configuration from site A:

```
metrocluster show
```

```

cluster_A::> metrocluster show

Cluster                Entry Name                State
-----
Local: cluster_A      Configuration state configured
Mode                   normal
AUSO Failure Domain  auso-on-cluster-
disaster
Remote: cluster_B    Configuration state configured
Mode                   normal
AUSO Failure Domain  auso-on-cluster-
disaster

```

b. Verify the configuration from site B:

```
metrocluster show
```

```

cluster_B::> metrocluster show

Cluster                Entry Name                State
-----
Local: cluster_B      Configuration state configured
Mode                   normal
AUSO Failure Domain  auso-on-cluster-
disaster
Remote: cluster_A    Configuration state configured
Mode                   normal
AUSO Failure Domain  auso-on-cluster-
disaster

```

Configuring in-order delivery or out-of-order delivery of frames on ONTAP software

You must configure either in-order delivery (IOD) or out-of-order delivery (OOD) of frames according to the fibre channel (FC) switch configuration.

About this task

If the FC switch is configured for IOD, then the ONTAP software must be configured for IOD. Similarly, if the FC switch is configured for OOD, then ONTAP must be configured for OOD.



You must reboot the controller to change the configuration.

Step

1. Configure ONTAP to operate either IOD or OOD of frames.
 - By default, IOD of frames is enabled in ONTAP. To check the configuration details:
 - a. Enter advanced mode:

```
set advanced
```

b. Verify the settings:

```
metrocluster interconnect adapter show
```

```
mcc4-b12_siteB::*> metrocluster interconnect adapter show
```

Node Port Number	Adapter Name	Adapter Type	Link Status	Is OOD Enabled?	IP Address
mcc4-b1 6a	fcvi_device_0	FC-VI	Up	false	17.0.1.2
mcc4-b1 6b	fcvi_device_1	FC-VI	Up	false	18.0.0.2
mcc4-b1 ib2a	mlx4_0	IB	Down	false	192.0.5.193
mcc4-b1 ib2b	mlx4_0	IB	Up	false	192.0.5.194
mcc4-b2 6a	fcvi_device_0	FC-VI	Up	false	17.0.2.2
mcc4-b2 6b	fcvi_device_1	FC-VI	Up	false	18.0.1.2
mcc4-b2 ib2a	mlx4_0	IB	Down	false	192.0.2.9
mcc4-b2 ib2b	mlx4_0	IB	Up	false	192.0.2.10

8 entries were displayed.

- The following steps must be performed on each node to configure OOD of frames:

a. Enter advanced mode:

```
set advanced
```

b. Verify the MetroCluster configuration settings:

```
metrocluster interconnect adapter show
```

```

mcc4-b12_siteB::*> metrocluster interconnect adapter show
                Adapter Link   Is OOD
Node           Adapter Name   Type   Status Enabled? IP Address
Port Number
-----
mcc4-b1        fcvi_device_0   FC-VI   Up    false  17.0.1.2
6a
mcc4-b1        fcvi_device_1   FC-VI   Up    false  18.0.0.2
6b
mcc4-b1        mlx4_0          IB      Down  false  192.0.5.193
ib2a
mcc4-b1        mlx4_0          IB      Up    false  192.0.5.194
ib2b
mcc4-b2        fcvi_device_0   FC-VI   Up    false  17.0.2.2
6a
mcc4-b2        fcvi_device_1   FC-VI   Up    false  18.0.1.2
6b
mcc4-b2        mlx4_0          IB      Down  false  192.0.2.9
ib2a
mcc4-b2        mlx4_0          IB      Up    false  192.0.2.10
ib2b
8 entries were displayed.

```

c. Enable OOD on node “mcc4-b1” and node “mcc4-b2”:

```

metrocluster interconnect adapter modify -node node_name -is-ood-enabled true

```

```

mcc4-b12_siteB::*> metrocluster interconnect adapter modify -node
mcc4-b1 -is-ood-enabled true
mcc4-b12_siteB::*> metrocluster interconnect adapter modify -node
mcc4-b2 -is-ood-enabled true

```

d. Reboot the controller by performing a high-availability (HA) takeover in both directions.

e. Verify the settings:

```

metrocluster interconnect adapter show

```

```

mcc4-b12_siteB::*> metrocluster interconnect adapter show
                        Adapter Link   Is OOD
Node                    Name      Type   Status Enabled? IP Address
Port Number
-----
mcc4-b1                 fcvi_device_0  FC-VI  Up     true   17.0.1.2
6a
mcc4-b1                 fcvi_device_1  FC-VI  Up     true   18.0.0.2
6b
mcc4-b1                 mlx4_0         IB     Down  false  192.0.5.193
ib2a
mcc4-b1                 mlx4_0         IB     Up    false  192.0.5.194
ib2b
mcc4-b2                 fcvi_device_0  FC-VI  Up     true   17.0.2.2
6a
mcc4-b2                 fcvi_device_1  FC-VI  Up     true   18.0.1.2
6b
mcc4-b2                 mlx4_0         IB     Down  false  192.0.2.9
ib2a
mcc4-b2                 mlx4_0         IB     Up    false  192.0.2.10
ib2b
8 entries were displayed.

```

Configuring SNMPv3 in a MetroCluster configuration

Before you begin

The authentication and privacy protocols on the switches and on the ONTAP system must be the same.

About this task

ONTAP currently supports AES-128 encryption.

Steps

1. Create an SNMP user for each switch from the controller prompt:

```
security login create
```

```

Controller_A_1::> security login create -user-or-group-name snmpv3user
-application snmp -authentication-method usm -role none -remote-switch
-ipaddress 10.10.10.10

```

2. Respond to the following prompts as required at your site:



For EngineID, press **ENTER** to assign the default value.

```
Enter the authoritative entity's EngineID [remote EngineID]:
```

```
Which authentication protocol do you want to choose (none, md5, sha, sha2-256) [none]: sha
```

```
Enter the authentication protocol password (minimum 8 characters long):
```

```
Enter the authentication protocol password again:
```

```
Which privacy protocol do you want to choose (none, des, aes128) [none]: aes128
```

```
Enter privacy protocol password (minimum 8 characters long):
```

```
Enter privacy protocol password again:
```



The same username can be added to different switches with different IP addresses.

3. Create an SNMP user for the rest of the switches.

The following example shows how to create a username for a switch with the IP address 10.10.10.11.

```
Controller_A_1::> security login create -user-or-group-name snmpv3user  
-application snmp -authentication-method usm -role none -remote-switch  
-ipaddress 10.  
10.10.11
```

4. Check that there is one login entry for each switch:

```
security login show
```

```
Controller_A_1::> security login show -user-or-group-name snmpv3user  
-fields remote-switch-ipaddress
```

```
vserver      user-or-group-name application authentication-method  
remote-switch-ipaddress
```

```
-----  
-----
```

```
node_A_1 SVM 1 snmpv3user      snmp      usm  
10.10.10.10
```

```
node_A_1 SVM 2 snmpv3user      snmp      usm  
10.10.10.11
```

```
node_A_1 SVM 3 snmpv3user      snmp      usm  
10.10.10.12
```

```
node_A_1 SVM 4 snmpv3user      snmp      usm  
10.10.10.13
```

```
4 entries were displayed.
```

5. Configure SNMPv3 on the switches from the switch prompt:

Brocade switches (FOS 9.0 and later)

```
snmpconfig --add snmpv3 -index <index> -user <user_name> -groupname <rw/ro>
-auth_proto <auth_protocol> -auth_passwd <auth_password> -priv_proto
<priv_protocol> -priv_passwd <priv_password>
```

Brocade switches (FOS 8.x and earlier)

```
snmpconfig --set snmpv3
```

The example shows how to configure a read-only user. You can adjust the RW users if needed. If you require RO access, after "User (ro):" specify the "snmpv3user".

```
Switch-A1:admin> snmpconfig --set snmpv3
SNMP Informs Enabled (true, t, false, f): [false] true
SNMPv3 user configuration(snmp user not configured in FOS user
database will have physical AD and admin role as the default):
User (rw): [snmpadmin1]
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [3]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [2]
Engine ID: [00:00:00:00:00:00:00:00]
User (ro): [snmpuser2] snmpv3user
Auth Protocol [MD5(1)/SHA(2)/noAuth(3)]: (1..3) [2]
Priv Protocol [DES(1)/noPriv(2)/AES128(3)/AES256(4)]: (2..2) [3]
```

Cisco switches

```
snmp-server user <user_name> auth [md5/sha/sha-256] <auth_password> priv
(aes-128) <priv_password>
```



You should also set passwords on unused accounts to secure them and use the best encryption available in your ONTAP release.

6. Configure encryption and passwords on the remaining switch users as required on your site.

Configuring MetroCluster components for health monitoring

You must perform some special configuration steps before monitoring the components in a MetroCluster configuration.



For improved security, NetApp recommends that you configure SNMPv2 or SNMPv3 to monitor the switch health.

About this task

These tasks apply only to systems with FC-to-SAS bridges.

Beginning in Fabric OS 9.0.1, SNMPv2 is not supported for health monitoring on Brocade switches, you must use SNMPv3 instead. If you are using SNMPv3, you must configure SNMPv3 in ONTAP before proceeding to the following section. For more details, see [Configuring SNMPv3 in a MetroCluster configuration](#).



- You should place bridges and a node management LIF in a dedicated network to avoid interference from other sources.
- If you use a dedicated network for health monitoring, then each node must have a node management LIF in that dedicated network.

NetApp only supports the following tools to monitor the components in a MetroCluster FC configuration:

- Brocade Network Advisor (BNA)
- Brocade SANnav
- Active IQ Config Advisor
- NetApp Health Monitoring (ONTAP)
- MetroCluster data collector (MC_DC)

Configuring the MetroCluster FC switches for health monitoring

In a fabric-attached MetroCluster configuration, you must perform some additional configuration steps to monitor the FC switches.



Beginning with ONTAP 9.8, the `storage switch` command is replaced with `system switch fibre-channel`. The following steps show the `storage switch` command, but if you are running ONTAP 9.8 or later, the `system switch fibre-channel` command is preferred.

Steps

1. Add a switch with an IP address to each MetroCluster node:

The command you run depends on whether you are using SNMPv2 or SNMPv3.

Add a switch using SNMPv3:

```
storage switch add -address <ip_address> -snmp-version SNMPv3 -snmp  
-community-or-username <SNMP_user_configured_on_the_switch>
```

Add a switch using SNMPv2:

```
storage switch add -address ipaddress
```

This command must be repeated on all four switches in the MetroCluster configuration.



Brocade 7840 FC switches and all alerts are supported in health monitoring, except `NoISLPresent_Alert`.

The following example shows the command to add a switch with IP address 10.10.10.10:

```
controller_A_1::> storage switch add -address 10.10.10.10
```

2. Verify that all switches are properly configured:

```
storage switch show
```

It might take up to 15 minutes to reflect all data due to the 15-minute polling interval.

The following example shows the command given to verify that the MetroCluster FC switches are configured:

```
controller_A_1::> storage switch show
Fabric          Switch Name      Vendor  Model          Switch WWN
Status
-----
-----
1000000533a9e7a6 brcd6505-fcs40  Brocade Brocade6505   1000000533a9e7a6
OK
1000000533a9e7a6 brcd6505-fcs42  Brocade Brocade6505   1000000533d3660a
OK
1000000533ed94d1 brcd6510-fcs44  Brocade Brocade6510   1000000533eda031
OK
1000000533ed94d1 brcd6510-fcs45  Brocade Brocade6510   1000000533ed94d1
OK
4 entries were displayed.

controller_A_1::>
```

If the worldwide name (WWN) of the switch is shown, the ONTAP health monitor can contact and monitor the FC switch.

Related information

[System administration](#)

Configuring FC-to-SAS bridges for health monitoring

In systems running ONTAP versions prior to 9.8, you must perform some special configuration steps to monitor the FC-to-SAS bridges in the MetroCluster configuration.

About this task

- Third-party SNMP monitoring tools are not supported for FibreBridge bridges.
- Beginning with ONTAP 9.8, FC-to-SAS bridges are monitored via in-band connections by default, and additional configuration is not required.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. From the ONTAP cluster prompt, add the bridge to health monitoring:
 - a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
---------------	---------

9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the "Status" column is "ok", and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```
controller_A_1::> storage bridge show

Bridge                Symbolic Name Is Monitored  Monitor Status
Vendor Model          Bridge WWN
-----
ATTO_10.10.20.10  atto01      true         ok           Atto
FibreBridge 7500N  20000010867038c0
ATTO_10.10.20.11  atto02      true         ok           Atto
FibreBridge 7500N  20000010867033c0
ATTO_10.10.20.12  atto03      true         ok           Atto
FibreBridge 7500N  20000010867030c0
ATTO_10.10.20.13  atto04      true         ok           Atto
FibreBridge 7500N  2000001086703b80

4 entries were displayed

controller_A_1::>
```

Checking the MetroCluster configuration

You can check that the components and relationships in the MetroCluster configuration are working correctly.

You should do a check after initial configuration and after making any changes to the MetroCluster configuration. You should also do a check before a negotiated (planned) switchover or a switchback operation.

About this task

If the `metrocluster check run` command is issued twice within a short time on either or both clusters, a conflict can occur and the command might not collect all data. Subsequent `metrocluster check show` commands, then will not show the expected output.

Steps

1. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```
cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results. To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"
```

```
cluster_A::> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

2. Display more detailed results from the most recent `metrocluster check run` command:

```
metrocluster check aggregate show
```

```
metrocluster check cluster show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```



The `metrocluster check show` commands show the results of the most recent `metrocluster check run` command. You should always run the `metrocluster check run` command prior to using the `metrocluster check show` commands so that the information displayed is current.

The following example shows the `metrocluster check aggregate show` command output for a

healthy four-node MetroCluster configuration:

```
cluster_A::> metrocluster check aggregate show
```

```
Last Checked On: 8/5/2014 00:42:58
```

Node	Aggregate	Check
Result		
-----	-----	-----
controller_A_1	controller_A_1_aggr0	mirroring-status
ok		disk-pool-allocation
ok		ownership-state
ok	controller_A_1_aggr1	mirroring-status
ok		disk-pool-allocation
ok		ownership-state
ok	controller_A_1_aggr2	mirroring-status
ok		disk-pool-allocation
ok		ownership-state
controller_A_2	controller_A_2_aggr0	mirroring-status
ok		disk-pool-allocation
ok		ownership-state
ok	controller_A_2_aggr1	mirroring-status
ok		disk-pool-allocation
ok		ownership-state

```

ok
        controller_A_2_aggr2
                                mirroring-status
ok
                                disk-pool-allocation
ok
                                ownership-state
ok

18 entries were displayed.

```

The following example shows the `metrocluster check cluster show` command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

```

Last Checked On: 9/13/2017 20:47:04

Cluster          Check          Result
-----
mccint-fas9000-0102
                negotiated-switchover-ready  not-applicable
                switchback-ready    not-applicable
                job-schedules      ok
                licenses          ok
                periodic-check-enabled ok
mccint-fas9000-0304
                negotiated-switchover-ready  not-applicable
                switchback-ready    not-applicable
                job-schedules      ok
                licenses          ok
                periodic-check-enabled ok

10 entries were displayed.

```

Related information

[Disk and aggregate management](#)

[Network and LIF management](#)

Checking for MetroCluster configuration errors with Config Advisor

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

Steps

1. Go to the Config Advisor download page and download the tool.

[NetApp Downloads: Config Advisor](#)

2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Verifying local HA operation

If you have a four-node MetroCluster configuration, you should verify the operation of the local HA pairs in the MetroCluster configuration. This is not required for two-node configurations.

About this task

Two-node MetroCluster configurations do not consist of local HA pairs and this task does not apply.

The examples in this task use standard naming conventions:

- cluster_A
 - controller_A_1
 - controller_A_2
- cluster_B
 - controller_B_1
 - controller_B_2

Steps

1. On cluster_A, perform a failover and giveback in both directions.
 - a. Confirm that storage failover is enabled:

```
storage failover show
```

The output should indicate that takeover is possible for both nodes:

```
cluster_A::> storage failover show
                                Takeover
Node           Partner           Possible State Description
-----
controller_A_1 controller_A_2 true      Connected to controller_A_2
controller_A_2 controller_A_1 true      Connected to controller_A_1
2 entries were displayed.
```

- b. Take over controller_A_2 from controller_A_1:

```
storage failover takeover controller_A_2
```

You can use the `storage failover show-takeover` command to monitor the progress of the takeover operation.

- c. Confirm that the takeover is complete:

```
storage failover show
```

The output should indicate that `controller_A_1` is in takeover state, meaning that it has taken over its HA partner:

```
cluster_A::> storage failover show
Node           Partner           Takeover
-----
controller_A_1 controller_A_2 false   In takeover
controller_A_2 controller_A_1 -      Unknown
2 entries were displayed.
```

- d. Give back controller_A_2:

```
storage failover giveback controller_A_2
```

You can use the `storage failover show-giveback` command to monitor the progress of the giveback operation.

- e. Confirm that storage failover has returned to a normal state:

```
storage failover show
```

The output should indicate that takeover is possible for both nodes:

```
cluster_A::> storage failover show
Node           Partner           Takeover
-----
controller_A_1 controller_A_2 true    Connected to controller_A_2
controller_A_2 controller_A_1 true    Connected to controller_A_1
2 entries were displayed.
```

- f. Repeat the previous substeps, this time taking over `controller_A_1` from `controller_A_2`.

2. Repeat the preceding steps on `cluster_B`.

Related information

[High-availability configuration](#)

Verifying switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

Step

1. Use the procedures for negotiated switchover, healing, and switchback that are mentioned in the [Recover from a disaster](#).

Protecting configuration backup files

You can provide additional protection for the cluster configuration backup files by specifying a remote URL (either HTTP or FTP) where the configuration backup files will be uploaded in addition to the default locations in the local cluster.

Step

1. Set the URL of the remote destination for the configuration backup files:

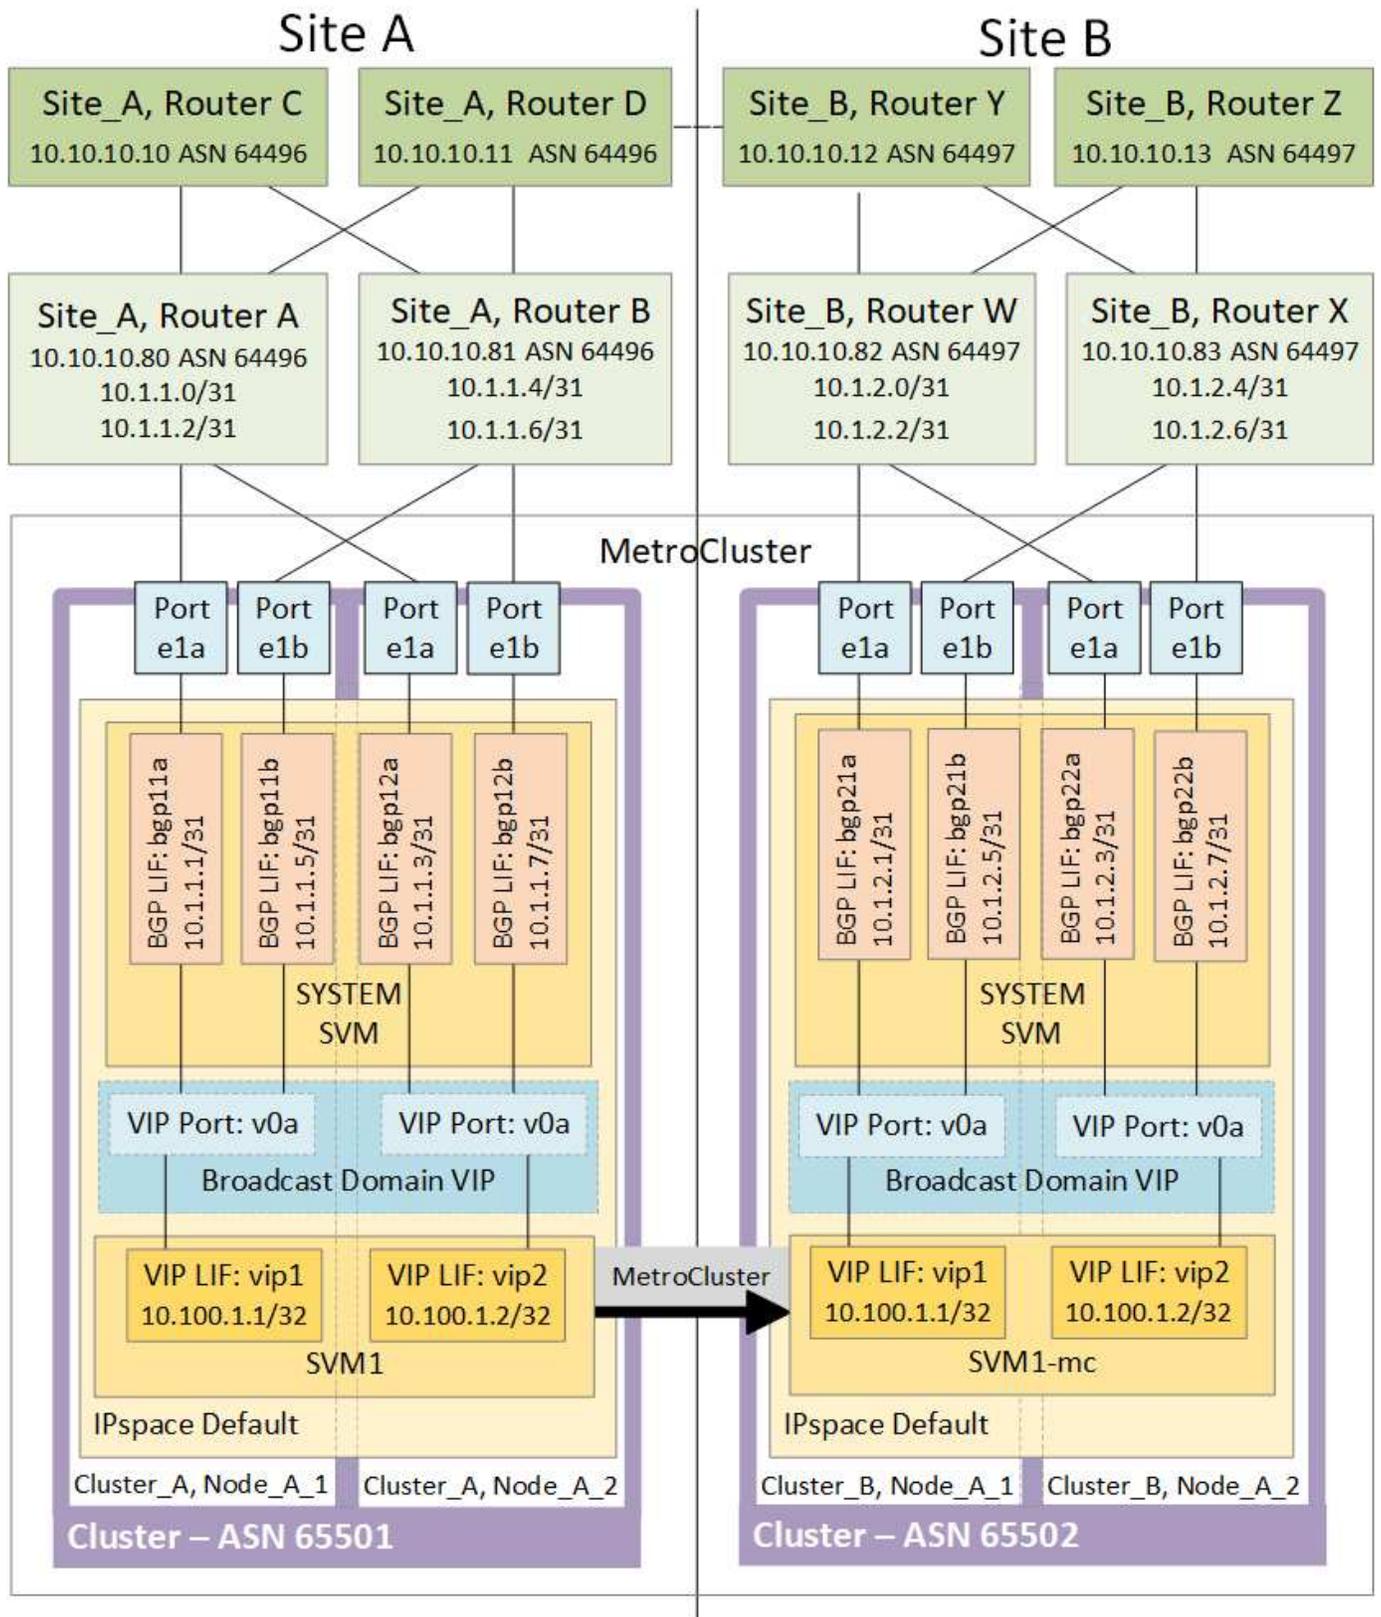
```
system configuration backup settings modify URL-of-destination
```

The [Cluster Management with the CLI](#) contains additional information under the section *Managing configuration backups*.

Considerations for using virtual IP and Border Gateway Protocol with a MetroCluster configuration

Beginning with ONTAP 9.5, ONTAP supports layer 3 connectivity using virtual IP (VIP) and Border Gateway Protocol (BGP). The combination VIP and BGP for redundancy in the front-end networking with the back-end MetroCluster redundancy provides a layer 3 disaster recovery solution.

Review the following guidelines and illustration when planning your layer 3 solution. For details on implementing VIP and BGP in ONTAP, refer to [Configure virtual IP LIFs](#).



ONTAP limitations

ONTAP does not automatically verify that all nodes on both sites of the MetroCluster configuration are configured with BGP peering.

ONTAP does not perform route aggregation but announces all individual virtual LIF IPs as unique host routes at all times.

ONTAP does not support true AnyCast — only a single node in the cluster presents a specific virtual LIF IP (but is accepted by all physical interfaces, regardless of whether they are BGP LIFs, provided the physical port is part of the correct IPspace). Different LIFs can migrate independently of each other to different hosting nodes.

Guidelines for using this Layer 3 solution with a MetroCluster configuration

You must configure your BGP and VIP correctly to provide the required redundancy.

Simpler deployment scenarios are preferred over more complex architectures (for example, a BGP peering router is reachable across an intermediate, non-BGP router). However, ONTAP does not enforce network design or topology restrictions.

VIP LIFs only cover the frontend/data network.

Depending on your version of ONTAP, you must configure BGP peering LIFs in the node SVM, not the system or data SVM. In ONTAP 9.8, the BGP LIFs are visible in the cluster (system) SVM and the node SVMs are no longer present.

Each data SVM requires the configuration of all potential first hop gateway addresses (typically, the BGP router peering IP address), so that the return data path is available if a LIF migration or MetroCluster failover occurs.

BGP LIFs are node specific, similar to intercluster LIFs — each node has a unique configuration, which does not need to be replicated to DR site nodes.

Once configured, the existence of the v0a (v0b and so on) continuously validates the connectivity, guaranteeing that a LIF migrate or failover succeeds (unlike L2, where a broken configuration is only visible after the outage).

A major architectural difference is that clients should no longer share the same IP subnet as the VIP of data SVMs. An L3 router with appropriate enterprise grade resiliency and redundancy features enabled (for example, VRRP/HSRP) should be on the path between storage and clients for the VIP to operate correctly.

The reliable update process of BGP allows for smoother LIF migrations because they are marginally faster and have a lower chance of interruption to some clients

You can configure BGP to detect some classes of network or switch misbehaviors faster than LACP, if configured accordingly.

External BGP (EBGP) uses different AS numbers between ONTAP node(s) and peering routers and is the preferred deployment to ease route aggregation and redistribution on the routers. Internal BGP (IBGP) and the use of route reflectors is not impossible but outside the scope of a straightforward VIP setup.

After deployment, you must check that the data SVM is accessible when the associated virtual LIF is migrated between all nodes on each site (including MetroCluster switchover) to verify the correct configuration of the static routes to the same data SVM.

VIP works for most IP-based protocols (NFS, SMB, iSCSI).

Testing the MetroCluster configuration

You can test failure scenarios to confirm the correct operation of the MetroCluster

configuration.

Verifying negotiated switchover

You can test the negotiated (planned) switchover operation to confirm uninterrupted data availability.

About this task

This test validates that data availability is not affected (except for SMB and Fibre Channel protocols) by switching the cluster over to the second data center.

This test should take about 30 minutes.

This procedure has the following expected results:

- The `metrocluster switchover` command will present a warning prompt.

If you respond `yes` to the prompt, the site the command is issued from will switch over the partner site.

For MetroCluster IP configurations:

- For ONTAP 9.4 and earlier:
 - Mirrored aggregates will become degraded after the negotiated switchover.
- For ONTAP 9.5 and later:
 - Mirrored aggregates will remain in normal state if the remote storage is accessible.
 - Mirrored aggregates will become degraded after the negotiated switchover if access to the remote storage is lost.
- For ONTAP 9.8 and later:
 - Unmirrored aggregates that are located at the disaster site will become unavailable if access to the remote storage is lost. This might lead to a controller outage.

Steps

1. Confirm that all nodes are in the configured state and normal mode:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show

Cluster                               Configuration State  Mode
-----
Local: cluster_A                       configured           normal
Remote: cluster_B                       configured           normal
```

2. Begin the switchover operation:

```
metrocluster switchover
```

```
cluster_A::> metrocluster switchover
Warning: negotiated switchover is about to start. It will stop all the
data Vservers on cluster "cluster_B" and
automatically re-start them on cluster "`cluster_A`". It will finally
gracefully shutdown cluster "cluster_B".
```

3. Confirm that the local cluster is in the configured state and switchover mode:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show

Cluster                               Configuration State      Mode
-----                               -
-----
Local: cluster_A                       configured                switchover
Remote: cluster_B                       not-reachable            -
      configured                    normal
```

4. Confirm that the switchover operation was successful:

```
metrocluster operation show
```

```
cluster_A::> metrocluster operation show
Operation: switchover
State: successful
Start Time: 2/6/2016 13:28:50
End Time: 2/6/2016 13:29:41
Errors: -
```

5. Use the `vserver show` and `network interface show` commands to verify that DR SVMs and LIFs have come online.

Verifying healing and manual switchback

You can test the healing and manual switchback operations to verify that data availability is not affected (except for SMB and Solaris FC configurations) by switching back the cluster to the original data center after a negotiated switchover.

About this task

This test should take about 30 minutes.

The expected result of this procedure is that services should be switched back to their home nodes.

Steps

1. Verify that healing is completed:

```
metrocluster node show
```

The following example shows the successful completion of the command:

```
cluster_A::> metrocluster node show
DR
Group Cluster Node          Configuration  DR
                               State            Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured    enabled    heal roots
completed
      cluster_B
      node_B_2      unreachable  -          switched over
42 entries were displayed.
```

2. Verify that all aggregates are mirrored:

```
storage aggregate show
```

The following example shows that all aggregates have a RAID Status of mirrored:

```

cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate Size      Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
data_cluster
      4.19TB      4.13TB   2% online    8 node_A_1  raid_dp,
mirrored,
normal

root_cluster
      715.5GB    212.7GB  70% online    1 node_A_1  raid4,
mirrored,
normal

cluster_B Switched Over Aggregates:
Aggregate Size      Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
data_cluster_B
      4.19TB      4.11TB   2% online    5 node_A_1  raid_dp,
mirrored,
normal

root_cluster_B    -          -      - unknown    - node_A_1  -

```

3. Boot the nodes from the disaster site.
4. Check the status of switchback recovery:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR
Group Cluster Node      Configuration  DR
State          Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured    enabled    heal roots
completed
      cluster_B
      node_B_2      configured    enabled    waiting for
switchback                                     recovery

2 entries were displayed.

```

5. Perform the switchback:

```
metrocluster switchback
```

```
cluster_A::> metrocluster switchback
[Job 938] Job succeeded: Switchback is successful. Verify switchback
```

6. Confirm the status of the nodes:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
DR
Group Cluster Node Configuration State DR Mirroring Mode
-----
-----
1 cluster_A
   node_A_1 configured enabled normal
  cluster_B
   node_B_2 configured enabled normal

2 entries were displayed.
```

7. Confirm the status:

```
metrocluster operation show
```

The output should show a successful state.

```
cluster_A::> metrocluster operation show
Operation: switchback
State: successful
Start Time: 2/6/2016 13:54:25
End Time: 2/6/2016 13:56:15
Errors: -
```

Loss of a single FC-to-SAS bridge

You can test the failure of a single FC-to-SAS bridge to make sure there is no single point of failure.

About this task

This test should take about 15 minutes.

This procedure has the following expected results:

- Errors should be generated as the bridge is switched off.
- No failover or loss of service should occur.
- Only one path from the controller module to the drives behind the bridge is available.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. Turn off the power supplies of the bridge.
2. Confirm that the bridge monitoring indicates an error:

```
storage bridge show
```

```
cluster_A::> storage bridge show
```

Monitor	Bridge	Symbolic Name	Vendor	Model	Bridge WWN	Is Monitored
ATTO_10.65.57.145		bridge_A_1	Atto	FibreBridge 6500N	200000108662d46c	true

```
error
```

3. Confirm that the drives behind the bridge are available with a single path:

```
storage disk error show
```

```

cluster_A::> storage disk error show
Disk              Error Type          Error Text
-----
-----
1.0.0             onedomain           1.0.0 (5000cca057729118): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.
1.0.1             onedomain           1.0.1 (5000cca057727364): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.
1.0.2             onedomain           1.0.2 (5000cca05772e9d4): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.
...
1.0.23            onedomain           1.0.23 (5000cca05772e9d4): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.

```

Verifying operation after power line disruption

You can test the MetroCluster configuration's response to the failure of a PDU.

About this task

The best practice is for each power supply unit (PSU) in a component to be connected to separate power supplies. If both PSUs are connected to the same power distribution unit (PDU) and an electrical disruption occurs, the site could down or a complete shelf might become unavailable. Failure of one power line is tested to confirm that there is no cabling mismatch that could cause a service disruption.

This test should take about 15 minutes.

This test requires turning off power to all left-hand PDUs and then all right-hand PDUs on all of the racks containing the MetroCluster components.

This procedure has the following expected results:

- Errors should be generated as the PDUs are disconnected.
- No failover or loss of service should occur.

Steps

1. Turn off the power of the PDUs on the left-hand side of the rack containing the MetroCluster components.
2. Monitor the result on the console:

```
system environment sensors show -state fault
```

```
storage shelf show -errors
```

```

cluster_A::> system environment sensors show -state fault

Node Sensor                State Value/Units Crit-Low Warn-Low Warn-Hi
Crit-Hi
-----
-----
node_A_1
    PSU1                    fault
                               PSU_OFF
    PSU1 Pwr In OK          fault
                               FAULT
node_A_2
    PSU1                    fault
                               PSU_OFF
    PSU1 Pwr In OK          fault
                               FAULT

4 entries were displayed.

cluster_A::> storage shelf show -errors
    Shelf Name: 1.1
    Shelf UID: 50:0a:09:80:03:6c:44:d5
    Serial Number: SHFHU1443000059

Error Type                Description
-----
Power                    Critical condition is detected in storage shelf
power supply unit "1". The unit might fail.Reconnect PSU1

```

3. Turn the power back on to the left-hand PDUs.
4. Make sure that ONTAP clears the error condition.
5. Repeat the previous steps with the right-hand PDUs.

Verifying operation after a switch fabric failure

You can disable a switch fabric to show that data availability is not affected by the loss.

About this task

This test should take about 15 minutes.

The expected result of this procedure is that disabling a fabric results in all cluster interconnect and disk traffic flowing to the other fabric.

In the examples shown, switch fabric 1 is disabled. This fabric consists of two switches, one at each MetroCluster site:

- FC_switch_A_1 on cluster_A

- FC_switch_B_1 on cluster_B

Steps

1. Disable connectivity to one of the two switch fabrics in the MetroCluster configuration:

- a. Disable the first switch in the fabric:

```
switchdisable
```

```
FC_switch_A_1::> switchdisable
```

- b. Disable the second switch in the fabric:

```
switchdisable
```

```
FC_switch_B_1::> switchdisable
```

2. Monitor the result on the console of the controller modules.

You can use the following commands to check the cluster nodes to make sure that all data is still being served. The command output shows missing paths to disks. This is expected.

- vserver show
- network interface show
- aggr show
- system node runnodename-command storage show disk -p
- storage disk error show

3. Reenable connectivity to one of the two switch fabrics in the MetroCluster configuration:

- a. Reenable the first switch in the fabric:

```
switchenable
```

```
FC_switch_A_1::> switchenable
```

- b. Reenable the second switch in the fabric:

```
switchenable
```

```
FC_switch_B_1::> switchenable
```

4. Wait at least 10 minutes and then repeat the above steps on the other switch fabric.

Verifying operation after loss of a single storage shelf

You can test the failure of a single storage shelf to verify that there is no single point of failure.

About this task

This procedure has the following expected results:

- An error message should be reported by the monitoring software.
- No failover or loss of service should occur.
- Mirror resynchronization starts automatically after the hardware failure is restored.

Steps

1. Check the storage failover status:

```
storage failover show
```

```
cluster_A::> storage failover show

Node           Partner           Possible State Description
-----
node_A_1       node_A_2           true      Connected to node_A_2
node_A_2       node_A_1           true      Connected to node_A_1
2 entries were displayed.
```

2. Check the aggregate status:

```
storage aggregate show
```

```
cluster_A::> storage aggregate show
```

```
cluster Aggregates:
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID

node_A_1data01_mirrored	4.15TB	3.40TB	18%	online	3	node_A_1	
raid_dp,							
mirrored,							
normal							
node_A_1root	707.7GB	34.29GB	95%	online	1	node_A_1	
raid_dp,							
mirrored,							
normal							
node_A_2_data01_mirrored	4.15TB	4.12TB	1%	online	2	node_A_2	
raid_dp,							
mirrored,							
normal							
node_A_2_data02_unmirrored	2.18TB	2.18TB	0%	online	1	node_A_2	
raid_dp,							
normal							
node_A_2_root	707.7GB	34.27GB	95%	online	1	node_A_2	
raid_dp,							
mirrored,							
normal							

3. Verify that all data SVMs and data volumes are online and serving data:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0,!MDV*
```

```
cluster_A::> vserver show -type data
```

```
cluster_A::> vserver show -type data
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume
Aggregate					
SVM1	data	sync-source		running	SVM1_root
node_A_1_data01_mirrored					
SVM2	data	sync-source		running	SVM2_root
node_A_2_data01_mirrored					

```
cluster_A::> network interface show -fields is-home false
```

```
There are no entries matching your query.
```

```
cluster_A::> volume show !vol0,!MDV*
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
SVM1					
	SVM1_root	node_A_1data01_mirrored	online	RW	10GB
9.50GB	5%				
SVM1					
	SVM1_data_vol	node_A_1data01_mirrored	online	RW	10GB
9.49GB	5%				
SVM2					
	SVM2_root	node_A_2_data01_mirrored	online	RW	10GB
9.49GB	5%				
SVM2					
	SVM2_data_vol	node_A_2_data02_unmirrored	online	RW	1GB
972.6MB	5%				

4. Identify a shelf in Pool 1 for node node_A_2 to power off to simulate a sudden hardware failure:

```
storage aggregate show -r -node node-name !*root
```

The shelf you select must contain drives that are part of a mirrored data aggregate.

In the following example, shelf ID 31 is selected to fail.

```
cluster_A::> storage aggregate show -r -node node_A_2 !*root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirrored) (block
checksums)
Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block
checksums)
```

Physical	Position	Disk	Pool	Type	RPM	Usable
Size	Status					Size
828.0GB (normal)	dparity	2.30.3	0	BSAS	7200	827.7GB
828.0GB (normal)	parity	2.30.4	0	BSAS	7200	827.7GB
828.0GB (normal)	data	2.30.6	0	BSAS	7200	827.7GB
828.0GB (normal)	data	2.30.8	0	BSAS	7200	827.7GB
828.0GB (normal)	data	2.30.5	0	BSAS	7200	827.7GB

```

Plex: /node_A_2_data01_mirrored/plex4 (online, normal, active, pool1)
RAID Group /node_A_2_data01_mirrored/plex4/rg0 (normal, block
checksums)
```

Physical	Position	Disk	Pool	Type	RPM	Usable
Size	Status					Size
828.0GB (normal)	dparity	1.31.7	1	BSAS	7200	827.7GB
828.0GB (normal)	parity	1.31.6	1	BSAS	7200	827.7GB
828.0GB (normal)	data	1.31.3	1	BSAS	7200	827.7GB

```

    data      1.31.4                1   BSAS      7200  827.7GB
828.0GB (normal)
    data      1.31.5                1   BSAS      7200  827.7GB
828.0GB (normal)

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block
checksums)
Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active,
pool0)
RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block
checksums)

                                                    Usable
Physical
  Position Disk                               Pool Type   RPM      Size
Size Status
-----
-----
    dparity  2.30.12                0   BSAS      7200  827.7GB
828.0GB (normal)
    parity   2.30.22                0   BSAS      7200  827.7GB
828.0GB (normal)
    data     2.30.21                0   BSAS      7200  827.7GB
828.0GB (normal)
    data     2.30.20                0   BSAS      7200  827.7GB
828.0GB (normal)
    data     2.30.14                0   BSAS      7200  827.7GB
828.0GB (normal)
15 entries were displayed.

```

5. Physically power off the shelf that you selected.

6. Check the aggregate status again:

```
storage aggregate show
```

```
storage aggregate show -r -node node_A_2 !*root
```

The aggregate with drives on the powered-off shelf should have a “degraded” RAID status, and drives on the affected plex should have a “failed” status, as shown in the following example:

```

cluster_A::> storage aggregate show
Aggregate      Size Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
node_A_1data01_mirrored
                4.15TB    3.40TB   18% online    3 node_A_1

```

```

raid_dp,

mirrored,

normal
node_A_1root
          707.7GB   34.29GB   95% online      1 node_A_1
raid_dp,

mirrored,

normal
node_A_2_data01_mirrored
          4.15TB    4.12TB    1% online      2 node_A_2
raid_dp,

mirror

degraded
node_A_2_data02_unmirrored
          2.18TB    2.18TB    0% online      1 node_A_2
raid_dp,

normal
node_A_2_root
          707.7GB   34.27GB   95% online      1 node_A_2
raid_dp,

mirror

```

```

degraded
cluster_A::> storage aggregate show -r -node node_A_2 !*root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirror degraded)
(block checksums)
Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block
checksums)

```

				Usable	
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	2.30.3	0	BSAS	7200	827.7GB
828.0GB (normal)					

```

    parity    2.30.4                0    BSAS    7200    827.7GB
828.0GB (normal)
    data      2.30.6                0    BSAS    7200    827.7GB
828.0GB (normal)
    data      2.30.8                0    BSAS    7200    827.7GB
828.0GB (normal)
    data      2.30.5                0    BSAS    7200    827.7GB
828.0GB (normal)

```

Plex: /node_A_2_data01_mirrored/plex4 (offline, failed, inactive, pool1)

RAID Group /node_A_2_data01_mirrored/plex4/rg0 (partial, none checksums)

					Usable
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	FAILED	-	-	-	827.7GB
- (failed)					
parity	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block checksums)

Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active, pool0)

RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block checksums)

					Usable
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	2.30.12	0	BSAS	7200	827.7GB
828.0GB (normal)					
parity	2.30.22	0	BSAS	7200	827.7GB
828.0GB (normal)					

```
data      2.30.21      0  BSAS  7200  827.7GB
828.0GB (normal)
data      2.30.20      0  BSAS  7200  827.7GB
828.0GB (normal)
data      2.30.14      0  BSAS  7200  827.7GB
828.0GB (normal)
```

15 entries were displayed.

7. Verify that the data is being served and that all volumes are still online:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0,!MDV*
```

```

cluster_A::> vserver show -type data

cluster_A::> vserver show -type data
Admin      Operational Root
Vserver    Type      Subtype    State      State      Volume
Aggregate
-----
-----
SVM1       data      sync-source  running    SVM1_root
node_A_1_data01_mirrored
SVM2       data      sync-source  running    SVM2_root
node_A_1_data01_mirrored

cluster_A::> network interface show -fields is-home false
There are no entries matching your query.

cluster_A::> volume show !vol0,!MDV*
Vserver    Volume      Aggregate    State      Type      Size
Available Used%
-----
-----
SVM1
          SVM1_root
                node_A_1data01_mirrored
                        online      RW      10GB
9.50GB    5%
SVM1
          SVM1_data_vol
                node_A_1data01_mirrored
                        online      RW      10GB
9.49GB    5%
SVM2
          SVM2_root
                node_A_1data01_mirrored
                        online      RW      10GB
9.49GB    5%
SVM2
          SVM2_data_vol
                node_A_2_data02_unmirrored
                        online      RW      1GB
972.6MB   5%

```

8. Physically power on the shelf.

Resynchronization starts automatically.

9. Verify that resynchronization has started:

```
storage aggregate show
```

The affected aggregate should have a “resyncing” RAID status, as shown in the following example:

```
cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
node_A_1_data01_mirrored
      4.15TB      3.40TB      18% online      3 node_A_1
raid_dp,
mirrored,
normal
node_A_1_root
      707.7GB      34.29GB      95% online      1 node_A_1
raid_dp,
mirrored,
normal
node_A_2_data01_mirrored
      4.15TB      4.12TB       1% online      2 node_A_2
raid_dp,
resyncing
node_A_2_data02_unmirrored
      2.18TB      2.18TB       0% online      1 node_A_2
raid_dp,
normal
node_A_2_root
      707.7GB      34.27GB      95% online      1 node_A_2
raid_dp,
resyncing
```

10. Monitor the aggregate to confirm that resynchronization is complete:

```
storage aggregate show
```

The affected aggregate should have a “normal” RAID status, as shown in the following example:

```

cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
node_A_1data01_mirrored
          4.15TB    3.40TB   18% online    3 node_A_1
raid_dp,

mirrored,

normal
node_A_1root
          707.7GB   34.29GB   95% online    1 node_A_1
raid_dp,

mirrored,

normal
node_A_2_data01_mirrored
          4.15TB    4.12TB    1% online    2 node_A_2
raid_dp,

normal
node_A_2_data02_unmirrored
          2.18TB    2.18TB    0% online    1 node_A_2
raid_dp,

normal
node_A_2_root
          707.7GB   34.27GB   95% online    1 node_A_2
raid_dp,

resyncing

```

Remove MetroCluster configurations

If you need to remove the MetroCluster configuration, contact technical support.

Contact NetApp technical support and reference the appropriate guide for your configuration from [How to remove nodes from a MetroCluster configuration - Resolution Guide](#).



You cannot reverse the MetroCluster unconfiguration. This process should only be done with the assistance of technical support. After removing the MetroCluster configuration, all disk connectivity and interconnects should be adjusted to be in a supported state.

How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

The Active IQ Unified Manager and ONTAP System Manager can be used for GUI management of the clusters and monitoring the configuration.

Each node has ONTAP System Manager pre-installed. To load System Manager, enter the cluster management LIF address as the URL in a web browser that has connectivity to the node.

You can also use Active IQ Unified Manager to monitor the MetroCluster configuration.

Related information

[Active IQ Unified Manager Documentation](#)

Synchronize the system time using NTP

Each cluster needs its own Network Time Protocol (NTP) server to synchronize the time between the nodes and their clients.

About this task

- You cannot modify the time zone settings for a failed node or the partner node after a takeover occurs.
- Each cluster in the MetroCluster FC configuration should have its own separate NTP server or servers used by the nodes, FC switches, and FC-to-SAS bridges at that MetroCluster site.
- If you are using the MetroCluster Tiebreaker software, it should also have its own separate NTP server.

Depending on your ONTAP version, you can configure the NTP from the **Cluster** or **Insights** tab in the System Manager UI.

Cluster

In System Manager, you can configure the NTP from the **Cluster** tab using two different options, depending on your ONTAP version:

ONTAP 9.8 or later:

Use the following steps to synchronize the NTP from the **Cluster** tab in ONTAP 9.8 or later.

Steps

1. Go to **Cluster > Overview**
2. Then select the  **More** option and select **Edit**.
3. In the **Edit Cluster Details** window, select the **+Add** option below NTP Servers.
4. Add the name, location, and specify the IP address of the time server.
5. Then, select **Save**.
6. Repeat the steps for any additional time servers.

ONTAP 9.11.1 or later:

Use the following steps to synchronize the NTP from the **Insights** window in the **Cluster** tab in ONTAP 9.11.1 or later.

Steps

1. Go to **Cluster > Overview**
2. Scroll down to the **Insights** window on the page, locate **Too few NTP servers are configured**, and then select **Fix It**.
3. Specify the IP address of the time server, and then select **Save**.
4. Repeat the previous step for any additional time servers.

Insights

In ONTAP 9.11.1 or later, you can also configure the NTP by using the **Insights** tab in System Manager:

Steps

1. Go to the **Insights** tab in the System Manager UI.
2. Scroll down to **Too few NTP servers are configured** and select **Fix It**.
3. Specify the IP address of the time server, and then select **Save**.
4. Repeat the previous step for any additional time servers.

Considerations when using ONTAP in a MetroCluster configuration

When using ONTAP in a MetroCluster configuration, you should be aware of certain considerations for licensing, peering to clusters outside the MetroCluster configuration, performing volume operations, NVFAIL operations, and other ONTAP operations.

Licensing considerations

- Both sites should be licensed for the same site-licensed features.
- All nodes should be licensed for the same node-locked features.

SnapMirror consideration

- SnapMirror SVM disaster recovery is only supported on MetroCluster configurations running versions of ONTAP 9.5 or later.

FlexCache support in a MetroCluster configuration

Beginning with ONTAP 9.7, FlexCache volumes are supported on MetroCluster configurations. You should be aware of requirements for manual repeer after switchover or switchback operations.

SVM repeer after switchover when FlexCache origin and cache are within the same MetroCluster site

After a negotiated or unplanned switchover, any SVM FlexCache peering relationship within the cluster must be manually configured.

For example, SVMs "vs1" (cache) and "vs2" (origin) are on site_A. These SVMs are peered.

After switchover, SVMs "vs1-mc" and "vs2-mc" are activated at the partner site (site_B). They must be manually repeer for FlexCache to work using the `vserver peer repeer` command.

SVM repeer after switchover or switchback when a FlexCache destination is on a third cluster and in disconnected mode

For FlexCache relationships to a cluster outside of the MetroCluster configuration, the peering must always be manually reconfigured after a switchover if the involved clusters are in disconnected mode during switchover.

For example:

- One end of the FlexCache (cache_1 on vs1) resides on MetroCluster site_A.
- The other end of the FlexCache (origin_1 on vs2) resides on site_C (not in the MetroCluster configuration).

When switchover is triggered, and if site_A and site_C are not connected, you must manually repeer the SVMs on site_B (the switchover cluster) and site_C using the `vserver peer repeer` command after the switchover.

When switchback is performed, you must again repeer the SVMs on site_A (the original cluster) and site_C.

Related information

[FlexCache volumes management with the CLI](#)

FabricPool support in MetroCluster configurations

Beginning with ONTAP 9.7, MetroCluster configurations support FabricPool storage tiers.

For general information on using FabricPools, see [Disk and tier \(aggregate\) management](#).

Considerations when using FabricPools

- The clusters must have FabricPool licenses with matching capacity limits.

- The clusters must have IPspaces with matching names.

This can be the default IPspace, or an IPspace that an administrator has created. This IPspace will be used for FabricPool object store configuration setups.

- For the selected IPspace, each cluster must have an intercluster LIF defined that can reach the external object store.
- SVM migration isn't supported with FabricPool when the source or destination is a MetroCluster cluster.

[Learn more about SVM data mobility.](#)

Configuring an aggregate for use in a mirrored FabricPool



Before you configure the aggregate, you must set up object stores as described in [Set up object stores for FabricPool in a MetroCluster configuration](#).

Steps

To configure an aggregate for use in a FabricPool:

1. Create the aggregate or select an existing aggregate.
2. Mirror the aggregate as a typical mirrored aggregate within the MetroCluster configuration.
3. Create the FabricPool mirror with the aggregate, as described in the [Disks and aggregates management](#)
 - a. Attach a primary object store.

This object store is physically closer to the cluster.

- b. Add a mirror object store.

This object store is physically further distant to the cluster than the primary object store.



It's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space may be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices may have a negative impact on performance.

FlexGroup support in MetroCluster configurations

Beginning with ONTAP 9.6 MetroCluster configurations support FlexGroup volumes.

Consistency group support in MetroCluster configurations

Beginning with ONTAP 9.11.1, [consistency groups](#) are supported in MetroCluster configurations.

Job schedules in a MetroCluster configuration

In ONTAP 9.3 and later, user-created job schedules are automatically replicated between clusters in a MetroCluster configuration. If you create, modify, or delete a job schedule on a cluster, the same schedule is automatically created on the partner cluster, using Configuration Replication Service (CRS).



System-created schedules are not replicated and you must manually perform the same operation on the partner cluster so that job schedules on both clusters are identical.

Cluster peering from the MetroCluster site to a third cluster

Because the peering configuration is not replicated, if you peer one of the clusters in the MetroCluster configuration to a third cluster outside of that configuration, you must also configure the peering on the partner MetroCluster cluster. This is so that peering can be maintained if a switchover occurs.

The non-MetroCluster cluster must be running ONTAP 8.3 or later. If not, peering is lost if a switchover occurs even if the peering has been configured on both MetroCluster partners.

LDAP client configuration replication in a MetroCluster configuration

An LDAP client configuration created on a storage virtual machine (SVM) on a local cluster is replicated to its partner data SVM on the remote cluster. For example, if the LDAP client configuration is created on the admin SVM on the local cluster, then it is replicated to all the admin data SVMs on the remote cluster. This MetroCluster feature is intentional so that the LDAP client configuration is active on all the partner SVMs on the remote cluster.

Networking and LIF creation guidelines for MetroCluster configurations

You should be aware of how LIFs are created and replicated in a MetroCluster configuration. You must also know about the requirement for consistency so that you can make proper decisions when configuring your network.

Related information

- [Network and LIF management](#)
- You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

[IPspace object replication and subnet configuration requirements](#)

- You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

[Requirements for LIF creation in a MetroCluster configuration](#)

- You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

[LIF replication and placement requirements and issues](#)

IPspace object replication and subnet configuration requirements

You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace replication

You must consider the following guidelines while replicating IPspace objects to the partner cluster:

- The IPspace names of the two sites must match.
- IPspace objects must be manually replicated to the partner cluster.

Any storage virtual machines (SVMs) that are created and assigned to an IPspace before the IPspace is replicated will not be replicated to the partner cluster.

Subnet configuration

You must consider the following guidelines while configuring subnets in a MetroCluster configuration:

- Both clusters of the MetroCluster configuration must have a subnet in the same IPspace with the same subnet name, subnet, broadcast domain, and gateway.
- The IP ranges of the two clusters must be different.

In the following example, the IP ranges are different:

```
cluster_A::> network subnet show

IPspace: Default
Subnet
Name      Subnet          Broadcast      Gateway      Avail/      Ranges
-----  -
subnet1   192.168.2.0/24   Default       192.168.2.1  10/10
192.168.2.11-192.168.2.20

cluster_B::> network subnet show
IPspace: Default
Subnet
Name      Subnet          Broadcast      Gateway      Avail/      Ranges
-----  -
subnet1   192.168.2.0/24   Default       192.168.2.1  10/10
192.168.2.21-192.168.2.30
```

IPv6 configuration

If IPv6 is configured on one site, IPv6 must be configured on the other site as well.

Related information

- You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

[Requirements for LIF creation in a MetroCluster configuration](#)

- You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

Requirements for LIF creation in a MetroCluster configuration

You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

You must consider the following guidelines when creating LIFs:

- Fibre Channel: You must use stretched VSAN or stretched fabrics
- IP/iSCSI: You must use layer 2 stretched network
- ARP broadcasts: You must enable ARP broadcasts between the two clusters
- Duplicate LIFs: You must not create multiple LIFs with the same IP address (duplicate LIFs) in an IPspace
- NFS and SAN configurations: You must use different storage virtual machines (SVMs) for both the unmirrored and mirrored aggregates
- You should create a subnet object before you create a LIF. A subnet object enables ONTAP to determine failover targets on the destination cluster because it has an associated broadcast domain.

Verify LIF creation

You can confirm the successful creation of a LIF in a MetroCluster configuration by running the `metrocluster check lif show` command. If you encounter any issues while creating the LIF, you can use the `metrocluster check lif repair-placement` command to fix the issues.

Related information

- You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

[IPspace object replication and subnet configuration requirements](#)

- You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

[LIF replication and placement requirements and issues](#)

LIF replication and placement requirements and issues

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

Replication of LIFs to the partner cluster

When you create a LIF on a cluster in a MetroCluster configuration, the LIF is replicated on the partner cluster. LIFs are not placed on a one-to-one name basis. For availability of LIFs after a switchover operation, the LIF placement process verifies that the ports are able to host the LIF based on reachability and port attribute checks.

The system must meet the following conditions to place the replicated LIFs on the partner cluster:

Condition	LIF type: FC	LIF type: IP/iSCSI
Node identification	ONTAP attempts to place the replicated LIF on the disaster recovery (DR) partner of the node on which it was created. If the DR partner is unavailable, the DR auxiliary partner is used for placement.	ONTAP attempts to place the replicated LIF on the DR partner of the node on which it was created. If the DR partner is unavailable, the DR auxiliary partner is used for placement.
Port identification	ONTAP identifies the connected FC target ports on the DR cluster.	<p>The ports on the DR cluster that are in the same IPspace as the source LIF are selected for a reachability check.</p> <p>If there are no ports in the DR cluster in the same IPspace, the LIF cannot be placed.</p> <p>All of the ports in the DR cluster that are already hosting a LIF in the same IPspace and subnet are automatically marked as reachable; and can be used for placement. These ports are not included in the reachability check.</p>
Reachability check	<p>Reachability is determined by checking for the connectivity of the source fabric WWN on the ports in the DR cluster.</p> <p>If the same fabric is not present at the DR site, the LIF is placed on a random port on the DR partner.</p>	<p>Reachability is determined by the response to an Address Resolution Protocol (ARP) broadcast from each previously identified port on the DR cluster to the source IP address of the LIF to be placed.</p> <p>For reachability checks to succeed, ARP broadcasts must be allowed between the two clusters.</p> <p>Each port that receives a response from the source LIF will be marked as possible for placement.</p>

<p>Port selection</p>	<p>ONTAP categorizes the ports based on attributes such as adapter type and speed, and then selects the ports with matching attributes.</p> <p>If no ports with matching attributes are found, the LIF is placed on a random connected port on the DR partner.</p>	<p>From the ports that are marked as reachable during the reachability check, ONTAP prefers ports that are in the broadcast domain that is associated with the subnet of the LIF.</p> <p>If there are no network ports available on the DR cluster that are in the broadcast domain that is associated with the subnet of the LIF, then ONTAP selects ports that have reachability to the source LIF.</p> <p>If there are no ports with reachability to the source LIF, a port is selected from the broadcast domain that is associated with the subnet of the source LIF, and if no such broadcast domain exists, a random port is selected.</p> <p>ONTAP categorizes the ports based on attributes such as adapter type, interface type, and speed, and then selects the ports with matching attributes.</p>
<p>LIF placement</p>	<p>From the reachable ports, ONTAP selects the least loaded port for placement.</p>	<p>From the selected ports, ONTAP selects the least loaded port for placement.</p>

Placement of replicated LIFs when the DR partner node is down

When an iSCSI or FC LIF is created on a node whose DR partner has been taken over, the replicated LIF is placed on the DR auxiliary partner node. After a subsequent giveback operation, the LIFs are not automatically moved to the DR partner. This can lead to LIFs being concentrated on a single node in the partner cluster. During a MetroCluster switchover operation, subsequent attempts to map LUNs belonging to the storage virtual machine (SVM) fail.

You should run the `metrocluster check lif show` command after a takeover operation or giveback operation to verify that the LIF placement is correct. If errors exist, you can run the `metrocluster check lif repair-placement` command to resolve the issues.

LIF placement errors

LIF placement errors that are displayed by the `metrocluster check lif show` command are retained after a switchover operation. If the `network interface modify`, `network interface rename`, or `network interface delete` command is issued for a LIF with a placement error, the error is removed and does not appear in the output of the `metrocluster check lif show` command.

LIF replication failure

You can also check whether LIF replication was successful by using the `metrocluster check lif show` command. An EMS message is displayed if LIF replication fails.

You can correct a replication failure by running the `metrocluster check lif repair-placement` command for any LIF that fails to find a correct port. You should resolve any LIF replication failures as soon as possible to verify the availability of LIF during a MetroCluster switchover operation.



Even if the source SVM is down, LIF placement might proceed normally if there is a LIF belonging to a different SVM in a port with the same IPspace and network in the destination SVM.

LIFs inaccessible after a switchover

If any change is made in the FC switch fabric to which the FC target ports of the source and DR nodes are connected, then the FC LIFs that are placed at the DR partner might become inaccessible to the hosts after a switchover operation.

You should run the `metrocluster check lif repair-placement` command on the source as well as the DR nodes after a change is made in the FC switch fabric to verify the host connectivity of LIFs. The changes in the switch fabric might result in LIFs getting placed in different target FC ports at the DR partner node.

Related information

- You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

[IPspace object replication and subnet configuration requirements](#)

- You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

[Requirements for LIF creation in a MetroCluster configuration](#)

Volume creation on a root aggregate

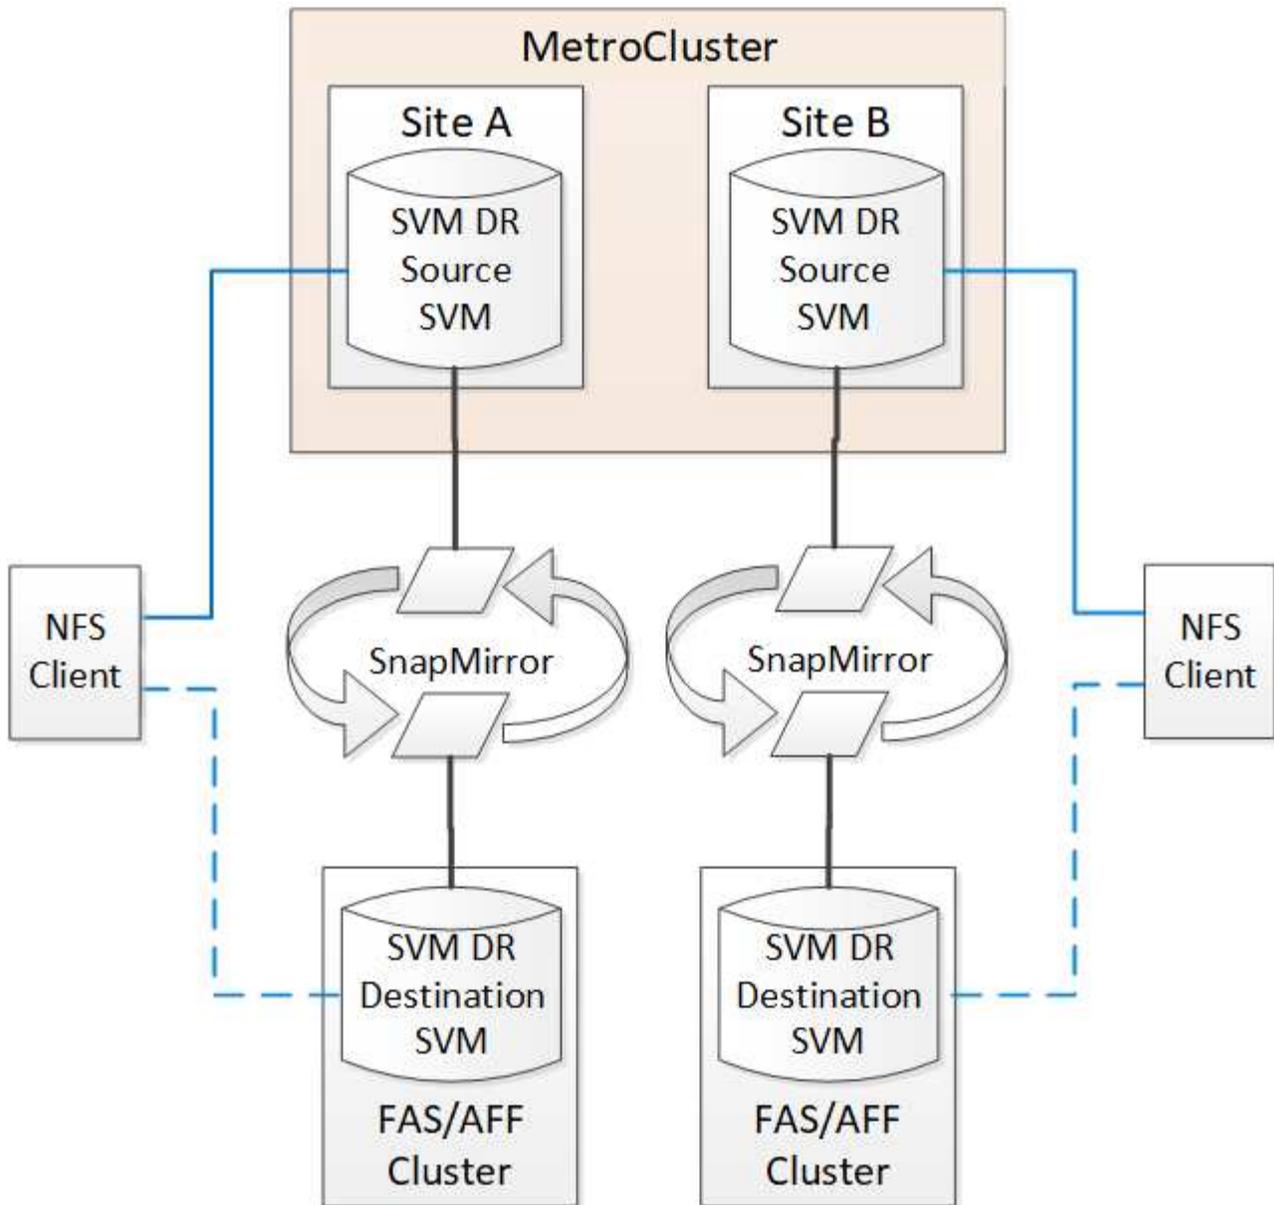
The system does not allow the creation of new volumes on the root aggregate (an aggregate with an HA policy of CFO) of a node in a MetroCluster configuration.

Because of this restriction, root aggregates cannot be added to an SVM using the `vserver add-aggregates` command.

SVM disaster recovery in a MetroCluster configuration

Beginning with ONTAP 9.5, active storage virtual machines (SVMs) in a MetroCluster configuration can be used as sources with the SnapMirror SVM disaster recovery feature. The destination SVM must be on the third cluster outside of the MetroCluster configuration.

Beginning with ONTAP 9.11.1, both sites within a MetroCluster configuration can be the source for an SVM DR relationship with a FAS or AFF destination cluster as shown in the following image.



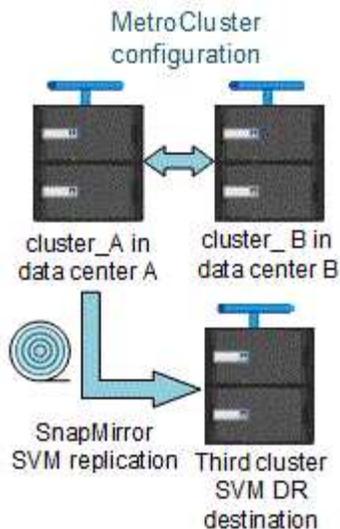
You should be aware of the following requirements and limitations of using SVMs with SnapMirror disaster recovery:

- Only an active SVM within a MetroCluster configuration can be the source of an SVM disaster recovery relationship.

A source can be a sync-source SVM before switchover or a sync-destination SVM after switchover.

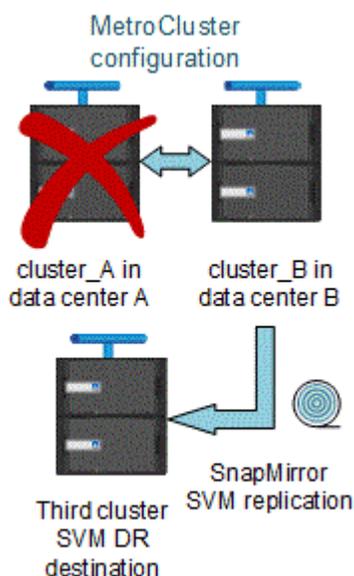
- When a MetroCluster configuration is in a steady state, the MetroCluster sync-destination SVM cannot be the source of an SVM disaster recovery relationship, since the volumes are not online.

The following image shows the SVM disaster recovery behavior in a steady state:



- When the sync-source SVM is the source of an SVM DR relationship, the source SVM DR relationship information is replicated to the MetroCluster partner.

This enables the SVM DR updates to continue after a switchover as shown in the following image:



- During the switchover and switchback processes, replication to the SVM DR destination might fail.

However, after the switchover or switchback process completes, the next SVM DR scheduled updates will succeed.

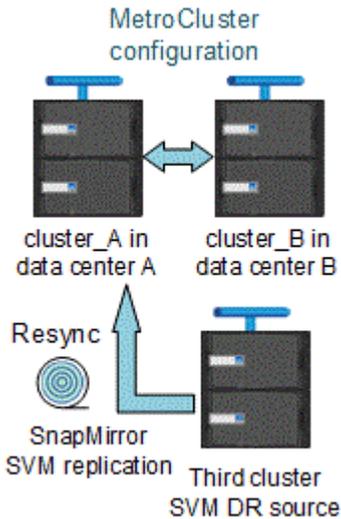
See the section “Replicating the SVM configuration” in the [Data Protection with the CLI](#) for details on configuring an SVM DR relationship.

SVM resynchronization at a disaster recovery site

During resynchronization, the storage virtual machines (SVMs) disaster recovery (DR) source on the MetroCluster configuration is restored from the destination SVM on the non-MetroCluster site.

During resynchronization, the source SVM (cluster_A) temporarily acts as a destination SVM as shown in the

following image:



If an unplanned switchover occurs during resynchronization

Unplanned switchovers that occur during the resynchronization will halt the resynchronization transfer. If an unplanned switchover occurs, the following conditions are true:

- The destination SVM on the MetroCluster site (which was a source SVM prior to resynchronization) remains as a destination SVM. The SVM at the partner cluster will continue to retain its subtype and remain inactive.
- The SnapMirror relationship must be re-created manually with the sync-destination SVM as the destination.
- The SnapMirror relationship does not appear in the SnapMirror show output after a switchover at the survivor site unless a SnapMirror create operation is executed.

Performing switchback after an unplanned switchover during resynchronization

To successfully perform the switchback process, the resynchronization relationship must be broken and deleted. Switchback is not permitted if there are any SnapMirror DR destination SVMs in the MetroCluster configuration or if the cluster has an SVM of subtype "dp-destination".

Output for the "storage aggregate plex show" command is indeterminate after a MetroCluster switchover

When you run the `storage aggregate plex show` command after a MetroCluster switchover, the status of plex0 of the switched over root aggregate is indeterminate and is displayed as "failed". During this time, the switched over root is not updated. The actual status of this plex can only be determined after the MetroCluster healing phase.

Modifying volumes to set the NVFAIL flag in case of switchover

You can modify a volume so that the NVFAIL flag is set on the volume in the event of a MetroCluster switchover. The NVFAIL flag causes the volume to be fenced off from any modification. This is required for volumes that need to be handled as if committed writes to the volume were lost after the switchover.

About this task



In ONTAP versions earlier than 9.0, the NVFAIL flag is used for each switchover. In ONTAP 9.0 and later versions, the unplanned switchover (USO) is used.

Step

1. Enable MetroCluster configuration to trigger NVFAIL on switchover by setting the `vol -dr-force -nvfail` parameter to "on":

```
vol modify -vserver vserver-name -volume volume-name -dr-force-nvfail on
```

Where to find additional information

You can learn more about MetroCluster configuration and operation.

MetroCluster and miscellaneous information

Information	Subject
ONTAP 9 Documentation	<ul style="list-style-type: none"> • All MetroCluster information
NetApp MetroCluster Solution Architecture and Design, TR-4705	<ul style="list-style-type: none"> • A technical overview of the MetroCluster FC configuration and operation. • Best practices for MetroCluster FC configuration.
MetroCluster IP Solution Architecture and Design, TR-4689	<ul style="list-style-type: none"> • A technical overview of the MetroCluster IP configuration and operation. • Best practices for a MetroCluster IP configuration.
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration: Differences among the ONTAP MetroCluster configurations	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the configuration • Configuring the MetroCluster in ONTAP
MetroCluster management and disaster recovery	<ul style="list-style-type: none"> • Understanding the MetroCluster configuration • Switchover, healing and switchback • Disaster recovery

<p>Maintain the MetroCluster components</p>	<ul style="list-style-type: none"> • Guidelines for maintenance in a MetroCluster FC configuration • Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches • Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Replacing hardware at a disaster site in a fabric-attached or stretch MetroCluster FC configuration • Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration. • Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.
<p>Transition from MetroCluster FC to MetroCluster IP</p> <p>MetroCluster Upgrade and Expansion Guide</p>	<ul style="list-style-type: none"> • Upgrading or refreshing a MetroCluster configuration • Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration • Expanding a MetroCluster configuration by adding additional nodes
<p>MetroCluster Tiebreaker Software installation and configuration</p>	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
<p>Active IQ Digital Advisor documentation</p> <p>NetApp Documentation: Product Guides and Resources</p>	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance
<p>Copy-based transition</p>	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems
<p>ONTAP concepts</p>	<ul style="list-style-type: none"> • How mirrored aggregates work

Install a MetroCluster IP configuration

MetroCluster IP installation workflow

To install your MetroCluster IP configuration, you must perform a number of procedures in the correct order.

- [Prepare for the installation and understand all requirements.](#)
- [Cable the components](#)
- [Configure the software](#)
- [Configure ONTAP mediator](#) (optional)
- [Test the configuration](#)

Prepare for the MetroCluster installation

ONTAP MetroCluster configurations support matrix

The various MetroCluster configurations have key differences in the required components.

In all configurations, each of the two MetroCluster sites are configured as an ONTAP cluster. In a two-node MetroCluster configuration, each node is configured as a single-node cluster.

Feature	IP configurations	Fabric attached configurations		Stretch configurations	
		Four- or eight-node	Two-node	Two-node bridge-attached	Two-node direct-attached
Number of controllers	Four or eight ¹	Four or eight	Two	Two	Two
Uses an FC switch storage fabric	No	Yes	Yes	No	No
Uses an IP switch storage fabric	Yes	No	No	No	No
Uses FC-to-SAS bridges	No	Yes	Yes	Yes	No
Uses direct-attached SAS storage	Yes (local attached only)	No	No	No	Yes

Supports ADP	Yes (beginning with ONTAP 9.4)	No	No	No	No
Supports local HA	Yes	Yes	No	No	No
Supports ONTAP automatic unplanned switchover (AUSO)	No	Yes	Yes	Yes	Yes
Supports unmirrored aggregates	Yes (beginning with ONTAP 9.8)	Yes	Yes	Yes	Yes
Supports ONTAP Mediator	Yes (beginning with ONTAP 9.7)	No	No	No	No
Supports MetroCluster Tiebreaker	Yes (not in combination with ONTAP Mediator)	Yes	Yes	Yes	Yes
Supports All SAN Arrays	Yes	Yes	Yes	Yes	Yes

Notes

- Review the following considerations for eight-node MetroCluster IP configurations:
 - Eight-node configurations are supported beginning with ONTAP 9.9.1.
 - Only NetApp-validated MetroCluster switches (ordered from NetApp) are supported.
 - Configurations using IP-routed (layer 3) backend connections are not supported.

Support for All SAN Array systems in MetroCluster configurations

Some of the All SAN Arrays (ASAs) are supported in MetroCluster configurations. In the MetroCluster documentation, the information for AFF models applies to the corresponding ASA system. For example, all cabling and other information for the AFF A400 system also applies to the ASA AFF A400 system.

Supported platform configurations are listed in the [NetApp Hardware Universe](#).

Differences between ONTAP Mediator and MetroCluster Tiebreaker

Beginning with ONTAP 9.7, you can use either the ONTAP Mediator-assisted automatic unplanned switchover (MAUSO) in the MetroCluster IP configuration or you can use the MetroCluster Tiebreaker software. It is not required to use the MAUSO or Tiebreaker software; however, if you choose to not use either of these services, you must [perform a](#)

[manual recovery](#) if a disaster occurs.

The different MetroCluster configurations perform automatic switchover under different circumstances:

- **MetroCluster FC configurations using the AUSO capability (not present in MetroCluster IP configurations)**

In these configurations, AUSO is initiated if controllers fail but the storage (and bridges, if present) remain operational.

- **MetroCluster IP configurations using ONTAP Mediator (ONTAP 9.7 and later)**

In these configurations, MAUSO is initiated in the same circumstances as AUSO, as described above, and also after a complete site failure (controllers, storage, and switches).

[Learn about how the ONTAP Mediator supports automatic unplanned switchover.](#)

- **MetroCluster IP or FC configurations using the Tiebreaker software in active mode**

In these configurations, the Tiebreaker initiates unplanned switchover after a complete site failure.

Before using the Tiebreaker software, review the [MetroCluster Tiebreaker Software installation and configuration](#)

Interoperability of ONTAP Mediator with other applications and appliances

You cannot use any third-party applications or appliances that can trigger a switchover in combination with ONTAP Mediator. In addition, monitoring a MetroCluster configuration with MetroCluster Tiebreaker software is not supported when using ONTAP Mediator.

Learn about remote storage and MetroCluster IP configurations

You should understand how the controllers access the remote storage and how the MetroCluster IP addresses work.

Access to remote storage in MetroCluster IP configurations

In MetroCluster IP configurations, the only way the local controllers can reach the remote storage pools is via the remote controllers. The IP switches are connected to the Ethernet ports on the controllers; they do not have direct connections to the disk shelves. If the remote controller is down, the local controllers cannot reach their remote storage pools.

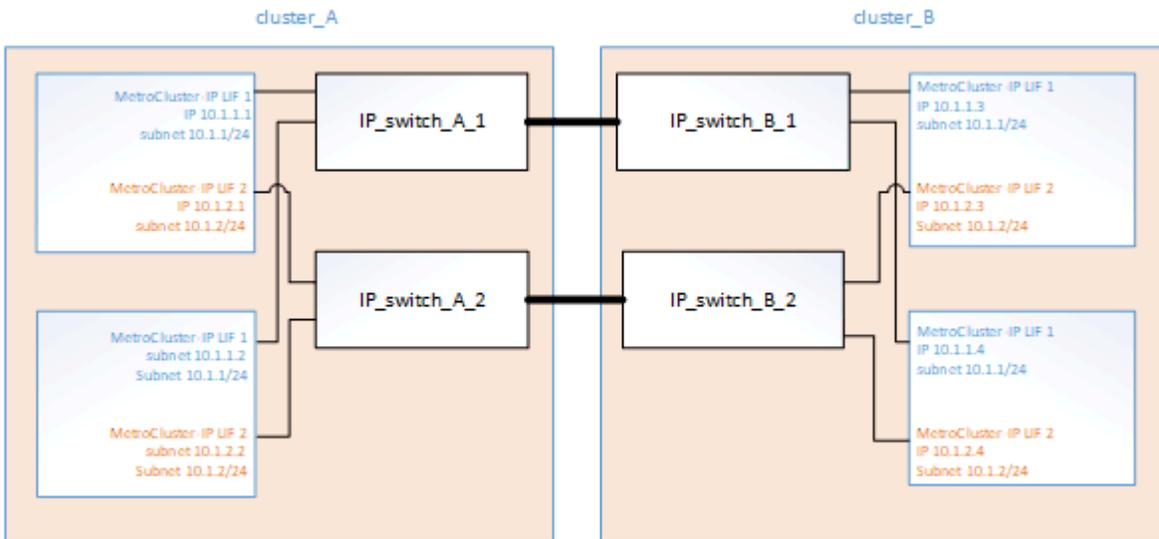
This is different than MetroCluster FC configurations, in which the remote storage pools are connected to the local controllers via the FC fabric or the SAS connections. The local controllers still have access to the remote storage even if the remote controllers are down.

MetroCluster IP addresses

You should be aware of how the MetroCluster IP addresses and interfaces are implemented in a MetroCluster IP configuration, as well as the associated requirements.

In a MetroCluster IP configuration, replication of storage and nonvolatile cache between the HA pairs and the DR partners is performed over high-bandwidth dedicated links in the MetroCluster IP fabric. iSCSI connections are used for storage replication. The IP switches are also used for all intra-cluster traffic within the local

clusters. The MetroCluster traffic is kept separate from the intra-cluster traffic by using separate IP subnets and VLANs. The MetroCluster IP fabric is distinct and different from the cluster peering network.



The MetroCluster IP configuration requires two IP addresses on each node that are reserved for the back-end MetroCluster IP fabric. The reserved IP addresses are assigned to MetroCluster IP logical interfaces (LIFs) during initial configuration, and have the following requirements:



You must choose the MetroCluster IP addresses carefully because you cannot change them after initial configuration.

- They must fall in a unique IP range.
They must not overlap with any IP space in the environment.
- They must reside in one of two IP subnets that separate them from all other traffic.

For example, the nodes might be configured with the following IP addresses:

Node	Interface	IP address	Subnet
node_A_1	MetroCluster IP interface 1	10.1.1.1	10.1.1/24
node_A_1	MetroCluster IP interface 2	10.1.2.1	10.1.2/24
node_A_2	MetroCluster IP interface 1	10.1.1.2	10.1.1/24
node_A_2	MetroCluster IP interface 2	10.1.2.2	10.1.2/24
node_B_1	MetroCluster IP interface 1	10.1.1.3	10.1.1/24

node_B_1	MetroCluster IP interface 2	10.1.2.3	10.1.2/24
node_B_2	MetroCluster IP interface 1	10.1.1.4	10.1.1/24
node_B_2	MetroCluster IP interface 2	10.1.2.4	10.1.2/24

Characteristics of MetroCluster IP interfaces

The MetroCluster IP interfaces are specific to MetroCluster IP configurations. They have different characteristics from other ONTAP interface types:

- They are created by the `metrocluster configuration-settings interface create` command as part the initial MetroCluster configuration.



Beginning with ONTAP 9.9.1, if you are using a layer 3 configuration, you must also specify the `-gateway` parameter when creating MetroCluster IP interfaces. Refer to [Considerations for layer 3 wide-area networks](#).

They are not created or modified by the network interface commands.

- They do not appear in the output of the `network interface show` command.
- They do not fail over, but remain associated with the port on which they were created.
- MetroCluster IP configurations use specific Ethernet ports (depending on the platform) for the MetroCluster IP interfaces.



Do not use 169.254.17.x or 169.254.18.x IP addresses when you create MetroCluster IP interfaces to avoid conflicts with system auto-generated interface IP addresses in the same range.

MetroCluster IP requirements for automatic drive assignment and ADP systems

Beginning with ONTAP 9.4, MetroCluster IP configurations support new installations using automatic disk assignment and ADP (Advanced Drive Partitioning).

You should be aware of the following when using ADP with MetroCluster IP configurations:

- ONTAP 9.4 and later is required to use ADP with MetroCluster IP configurations on AFF and ASA systems.
- ADPv2 is supported in MetroCluster IP configurations.
- The root aggregate must be located in Partition 3 for all nodes on both sites.
- Partitioning and disk assignment are performed automatically during the initial configuration of the MetroCluster sites.
- Pool 0 disk assignments are done at the factory.
- The unmirrored root is created at the factory.

- Data partition assignment is done at the customer site during the setup procedure.
- In most cases, drive assignment and partitioning is done automatically during the setup procedures.
- A disk and all of its partitions must be owned by nodes in the same high-availability (HA) pair. Partition or drive ownership within a single drive cannot be mixed between the local HA pair and the disaster recovery (DR) partner or DR auxiliary partner.

Example of a supported configuration:

Drive/Partition	Owner
Drive:	ClusterA-Node01
Partition 1:	ClusterA-Node01
Partition 2:	ClusterA-Node02
Partition 3:	ClusterA-Node01



When upgrading from ONTAP 9.4 to 9.5, the system recognizes the existing disk assignments.

Automatic partitioning

ADP is performed automatically during initial configuration of the system.



Beginning with ONTAP 9.5, automatic assignment of disks must be enabled with the `storage disk option modify -autoassign on` command.

You must set the `ha-config` state to `mccip` before automatic provisioning to make sure that the correct partition sizes are selected to allow for appropriate root volume size. For more information, see [Verifying the ha-config state of components](#).

A maximum of 96 drives can be automatically partitioned during installation. You can add extra drives after the initial installation.



If you are using internal and external drives, you first initialize the MetroCluster with only the internal drives using ADP. You then manually connect the external shelf after you complete your installation or setup task.

You must ensure that the internal shelves have the recommended minimum number of drives as outlined in [ADP and disk assignment differences by system](#).

For both the internal and external drives, you must populate the partially full shelves as described in [How to populate partially-full shelves](#).

How shelf-by-shelf automatic assignment works

If there are four external shelves per site, each shelf is assigned to a different node and different pool, as shown in the following example:

- All of the disks on `site_A-shelf_1` are automatically assigned to pool 0 of `node_A_1`
- All of the disks on `site_A-shelf_3` are automatically assigned to pool 0 of `node_A_2`

- All of the disks on site_B-shelf_1 are automatically assigned to pool 0 of node_B_1
- All of the disks on site_B-shelf_3 are automatically assigned to pool 0 of node_B_2
- All of the disks on site_B-shelf_2 are automatically assigned to pool 1 of node_A_1
- All of the disks on site_B-shelf_4 are automatically assigned to pool 1 of node_A_2
- All of the disks on site_A-shelf_2 are automatically assigned to pool 1 of node_B_1
- All of the disks on site_A-shelf_4 are automatically assigned to pool 1 of node_B_2

How to populate partially-full shelves

If your configuration is using shelves that are not fully populated (have empty drive bays) you must distribute the drives evenly throughout the shelf, depending on the disk assignment policy. The disk assignment policy depends on how many shelves are at each MetroCluster site.

If you are using a single shelf at each site (or just the internal shelf on an AFF A800 system), disks are assigned using a quarter-shelf policy. If the shelf is not fully populated, install the drives equally on all quarters.

The following table shows an example of how to place 24 disks in a 48 drive internal shelf. The ownership for the drives is also shown.

The 48 drive bays are divided into four quarters:	Install six drives in the first six bays in each quarter...
Quarter 1: Bays 0-11	Bays 0-5
Quarter 2: Bays 12-23	Bays 12-17
Quarter 3: Bays 24-35	Bays 24-29
Quarter 4: Bays 36-47	Bays 36-41

The following table shows an example of how to place 16 disks in a 24 drive internal shelf.

The 24 drive bays are divided into four quarters:	Install four drives in the first four bays in each quarter...
Quarter 1: Bays 0-5	Bays 0-3
Quarter 2: Bays 6-11	Bays 6-9
Quarter 3: Bays 12-17	Bays 12-15
Quarter 4: Bays 18-23	Bays 18-21

If you are using two external shelves at each site, disks are assigned using a half-shelf policy. If the shelves are not fully populated, install the drives equally from either end of the shelf.

For example, if you are installing 12 drives in a 24-drive shelf, install drives in bays 0-5 and 18-23.

Manual drive assignment (ONTAP 9.5)

In ONTAP 9.5, manual drive assignment is required on systems with the following shelf configurations:

- Three external shelves per site.

Two shelves are assigned automatically using a half-shelf assignment policy, but the third shelf must be assigned manually.

- More than four shelves per site and the total number of external shelves is not a multiple of four.

Extra shelves above the nearest multiple of four are left unassigned and the drives must be assigned manually. For example, if there are five external shelves at the site, shelf five must be assigned manually.

You only need to manually assign a single drive on each unassigned shelf. The rest of the drives on the shelf are then automatically assigned.

Manual drive assignment (ONTAP 9.4)

In ONTAP 9.4, manual drive assignment is required on systems with the following shelf configurations:

- Fewer than four external shelves per site.

The drives must be assigned manually to ensure symmetrical assignment of the drives, with each pool having an equal number of drives.

- More than four external shelves per site and the total number of external shelves is not a multiple of four.

Extra shelves above the nearest multiple of four are left unassigned and the drives must be assigned manually.

When manually assigning drives, you should assign disks symmetrically, with an equal number of drives assigned to each pool. For example, if the configuration has two storage shelves at each site, you would one shelf to the local HA pair and one shelf to the remote HA pair:

- Assign half of the disks on site_A-shelf_1 to pool 0 of node_A_1.
- Assign half of the disks on site_A-shelf_1 to pool 0 of node_A_2.
- Assign half of the disks on site_A-shelf_2 to pool 1 of node_B_1.
- Assign half of the disks on site_A-shelf_2 to pool 1 of node_B_2.
- Assign half of the disks on site_B-shelf_1 to pool 0 of node_B_1.
- Assign half of the disks on site_B-shelf_1 to pool 0 of node_B_2.
- Assign half of the disks on site_B-shelf_2 to pool 1 of node_A_1.
- Assign half of the disks on site_B-shelf_2 to pool 1 of node_A_2.

Adding shelves to an existing configuration

Automatic drive assignment supports the symmetrical addition of shelves to an existing configuration.

When new shelves are added, the system applies the same assignment policy to newly added shelves. For example, with a single shelf per site, if an additional shelf is added, the systems applies the quarter-shelf assignment rules to the new shelf.

Related information

[Required MetroCluster IP components and naming conventions](#)

[Disk and aggregate management](#)

ADP and disk assignment differences by system in MetroCluster IP configurations

The operation of Advanced Drive Partitioning (ADP) and automatic disk assignment in MetroCluster IP configurations varies depending on the system model.



In systems using ADP, aggregates are created using partitions in which each drive is partitioned in to P1, P2 and P3 partitions. The root aggregate is created using P3 partitions.

Review the following requirements before using the tables:

- You must meet the MetroCluster limits for the maximum number of supported drives and other guidelines. Refer to the [NetApp Hardware Universe](#).
- If you are reusing an external disk shelf, verify that disk ownership on the external disk shelf has been removed before you attach it to the controller. Refer to [Remove ONTAP ownership from a disk](#).

ADP and disk assignment on AFF A320 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	48 drives	The drives on each external shelf are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	One shelf is used by the local HA pair. The second shelf is used by the remote HA pair. Partitions on each shelf are used to create the root aggregate. Each of the two plexes in the root aggregate includes the following partitions: <ul style="list-style-type: none">• Eight partitions for data• Two parity partitions• Two spare partitions

Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.	Each of the two plexes in the root aggregate includes the following partitions: <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partition
-------------------------------------	-----------	---	---

ADP and disk assignment on AFF A150, ASA A150, and AFF A220 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	Internal drives only	<p>The internal drives are divided into four equal groups. Each group is automatically assigned to a separate pool and each pool is assigned to a separate controller in the configuration.</p> <p>Note: Half of the internal drives remain unassigned before MetroCluster is configured.</p>	<p>Two quarters are used by the local HA pair. The other two quarters are used by the remote HA pair.</p> <p>The root aggregate includes the following partitions in each plex:</p> <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partition

<p>Minimum supported drives (per site)</p>	<p>16 internal drives</p>	<p>The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.</p> <p>Two quarters on a shelf can have the same pool. The pool is chosen based on the node that owns the quarter:</p> <ul style="list-style-type: none"> • If owned by the local node, pool0 is used. • If owned by the remote node, pool1 is used. <p>For example: a shelf with quarters Q1 through Q4 can have following assignments:</p> <ul style="list-style-type: none"> • Q1: node_A_1 pool0 • Q2: node_A_2 pool0 • Q3: node_B_1 pool1 • Q4:node_B_2 pool1 <p>Note: Half of the internal drives remain unassigned before MetroCluster is configured.</p>	<p>Each of the two plexes in the root aggregate includes the following partitions:</p> <ul style="list-style-type: none"> • Two partitions for data • Two parity partitions • No spares
--	---------------------------	--	--

ADP and disk assignment on AFF A250, AFF C250, ASA A250, ASA C250, FAS500f, AFF A20, AFF A30, and AFF C30 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
-----------	-----------------	------------------------	-------------------------------

Minimum recommended drives (per site)	48 drives (external drives only, no internal drives)	The drives on each external shelf are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	<p>One shelf is used by the local HA pair. The second shelf is used by the remote HA pair.</p> <p>Partitions on each shelf are used to create the root aggregate. The root aggregate includes the following partitions in each plex:</p> <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions
	48 drives (external and internal drives)	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool. The drives on the external shelves are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	<p>Each of the two plexes in the root aggregate includes:</p> <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	16 internal drives	The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.	<p>Each of the two plexes in the root aggregate includes the following partitions:</p> <ul style="list-style-type: none"> • Two partitions for data • Two parity partitions • No spare partitions

ADP and disk assignment on AFF A50 and AFF C60 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
------------------	------------------------	-------------------------------	--------------------------------------

Minimum recommended drives (per site)	48 drives (external drives only, no internal drives)	The drives on each external shelf are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	<p>The local HA pair uses one shelf. The remote HA pair uses the second shelf.</p> <p>Partitions on each shelf are used to create the root aggregate. The root aggregate includes the following partitions in each plex:</p> <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions
	48 drives (external and internal drives)	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool. The drives on the external shelves are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	<p>Each of the two plexes in the root aggregate includes:</p> <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	24 internal drives	The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.	<p>Each of the two plexes in the root aggregate includes the following partitions:</p> <ul style="list-style-type: none"> • Two partitions for data • Two parity partitions • No spare partitions

ADP and disk assignment on AFF A300 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
------------------	------------------------	-------------------------------	--------------------------------------

Minimum recommended drives (per site)	48 drives	The drives on each external shelf are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	<p>One shelf is used by the local HA pair. The second shelf is used by the remote HA pair.</p> <p>Partitions on each shelf are used to create the root aggregate. The root aggregate includes the following partitions in each plex:</p> <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.	<p>Each of the two plexes in the root aggregate includes the following partitions:</p> <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partition

ADP and disk assignment on AFF C400, AFF A400, ASA C400, and ASA A400 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	96 drives	Drives are automatically assigned on a shelf-by-shelf basis.	<p>Each of the two plexes in the root aggregate includes:</p> <ul style="list-style-type: none"> • 20 partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	<p>Each of the two plexes in the root aggregate includes:</p> <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partition

ADP and disk assignment on AFF A700 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	96 drives	Drives are automatically assigned on a shelf-by-shelf basis.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • 20 partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partition

ADP and disk assignment on AFF C800, ASA C800, ASA A800, and AFF A800 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root aggregate
Minimum recommended drives (per site)	Internal drives and 96 external drives	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool. The drives on the external shelves are automatically assigned on a shelf-by-shelf basis, with all of the drives on each shelf assigned to one of the four nodes in the MetroCluster configuration.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions <p>Note: The root aggregate is created with 12 root partitions on the internal shelf.</p>

Minimum supported drives (per site)	24 internal drives	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partitions <p>Note: The root aggregate is created with 12 root partitions on the internal shelf.</p>
-------------------------------------	--------------------	---	--

ADP and disk assignment on AFF A70, AFF A90, and AFF C80 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root aggregate
Minimum recommended drives (per site)	Internal drives and 96 external drives	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool. The drives on the external shelves are automatically assigned on a shelf-by-shelf basis, with all of the drives on each shelf assigned to one of the four nodes in the MetroCluster configuration.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • Eight partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	24 internal drives	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partitions

ADP and disk assignment on AFF A900, ASA A900, and AFF A1K systems

Guideline	Shelves per site	Drive assignment rules	ADP layout for root partition
-----------	------------------	------------------------	-------------------------------

Minimum recommended drives (per site)	96 drives	Drives are automatically assigned on a shelf-by-shelf basis.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • 20 partitions for data • Two parity partitions • Two spare partitions
Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	Each of the two plexes in the root aggregate includes: <ul style="list-style-type: none"> • Three partitions for data • Two parity partitions • One spare partition

Disk assignment on FAS2750 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	24 internal drives and 24 external drives	The internal and external shelves are divided into two equal halves. Each half is automatically assigned to different pool	Not applicable
Minimum supported drives (per site) (active/passive HA configuration)	Internal drives only	Manual assignment required	Not applicable

Disk assignment on FAS8200 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	48 drives	The drives on the external shelves are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	Not applicable

Minimum supported drives (per site) (active/passive HA configuration)	24 drives	Manual assignment required.	Not applicable
---	-----------	-----------------------------	----------------

Disk assignment on FAS500f systems

The same disk assignment guidelines and rules for AFF C250 and AFF A250 systems apply to FAS500f systems. For disk assignment on FAS500f systems, refer to the [ADP and disk assignment on AFF A250, AFF C250, ASA A250, ASA C250, FAS500f, AFF A20, AFF A30, and AFF C30 systems](#) table.

Disk assignment on FAS9000, FAS9500, FAS70, and FAS90 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
Minimum recommended drives (per site)	96 drives	Drives are automatically assigned on a shelf-by-shelf basis.	Not applicable
Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	Not applicable

Disk assignment on FAS50 systems

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
------------------	------------------------	-------------------------------	--------------------------------------

Minimum recommended drives (per site)	48 drives (external drives only, no internal drives)	The drives on each external shelf are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	Not applicable
	48 drives (external and internal drives)	The internal partitions are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool. The drives on the external shelves are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.	Not applicable
Minimum supported drives (per site)	24 drives	The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.	Not applicable

Disk assignment on FAS50 systems with Flash Cache

Beginning with ONTAP 9.18.1, Flash Cache is supported on FAS50 systems for MetroCluster IP configurations.



- You cannot have both data aggregates and the root aggregate with Flash Cache drives on the internal shelf.
- Slots 0 and 23 are used for Flash Cache drives.
- If you are reusing an external disk shelf, verify that disk ownership on the external disk shelf has been removed before you attach it to the controller. Refer to [Remove ONTAP ownership from a disk](#).

Guideline	Drives per site	Drive assignment rules	ADP layout for root partition
-----------	-----------------	------------------------	-------------------------------

Minimum recommended drives (per site)	48 drives (external drives only, no internal drives)	The drives on each external shelf are divided into two equal groups (halves). Each half-shelf is automatically assigned to a separate pool.	Not applicable
	36 drives (12 internal drives and 24 external drives – with the data aggregates on the external shelf and the root aggregate on the internal shelf)	<p>The internal drives are divided into four equal groups (quarters). Each quarter is automatically assigned to a separate pool. The drives on the external shelves are divided into four equal groups (quarters). Each quarter-shelf is automatically assigned to a separate pool.</p> <p>Notes:</p> <ul style="list-style-type: none"> • If you are using internal and external drives, you first need to install ONTAP and configure the local cluster with only the internal drives. You then manually connect the external shelf after you complete your installation or setup task. • A minimum of 12 internal drives is required for the root aggregate. You should place the root internal drives from slot 1. For example, for 12 internal drives, use slots 1-3, 6-8, 12-14, and 18-20. 	Not applicable
Minimum supported drives (per site)	24 external drives	The drives are divided into four equal groups. Each quarter-shelf is automatically assigned to a separate pool.	Not applicable

Requirements for cluster peering in MetroCluster IP configurations

Each MetroCluster site is configured as a peer to its partner site. You must be familiar with the prerequisites and guidelines for configuring the peering relationships. This is important when deciding on whether to use shared or dedicated ports for those relationships.

Related information

[Cluster and SVM peering express configuration](#)

Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that connectivity between port, IP address, subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node.

For example, in a four-node cluster, the subnet used for intercluster communication must have four available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.



ONTAP 9 enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

- All ports used to communicate with a given remote cluster must be in the same IPspace.

You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.

- The broadcast domain used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).

- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
- TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
- Bidirectional HTTPS between the intercluster LIFs

The default intercluster firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Considerations when using dedicated ports

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether using a dedicated port is the best intercluster network solution:

- If the amount of available WAN bandwidth is similar to that of the LAN ports, and the replication interval is such that replication occurs while regular client activity exists, then you should dedicate Ethernet ports for intercluster replication to avoid contention between replication and the data protocols.
- If the network utilization generated by the data protocols (CIFS, NFS, and iSCSI) is such that the network utilization is above 50 percent, then dedicate ports for replication to allow for nondegraded performance if a node failover occurs.
- When physical 10 GbE or faster ports are used for data and replication, you can create VLAN ports for replication and dedicate the logical ports for intercluster replication.

The bandwidth of the port is shared between all VLANs and the base port.

- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.

Considerations when sharing data ports

When determining whether sharing a data port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether sharing data ports is the best intercluster connectivity solution:

- For a high-speed network, such as a 40-Gigabit Ethernet (40-GbE) network, a sufficient amount of local LAN bandwidth might be available to perform replication on the same 40-GbE ports that are used for data access.

In many cases, the available WAN bandwidth is far less than the 10 GbE LAN bandwidth.

- All nodes in the cluster might have to replicate data and share the available WAN bandwidth, making data port sharing more acceptable.
- Sharing ports for data and replication eliminates the extra port counts required to dedicate ports for replication.
- The maximum transmission unit (MTU) size of the replication network will be the same size as that used on the data network.
- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.
- When data ports for intercluster replication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node to control the specific data port that is used for replication.

ISL requirements

Inter-Switch Link requirements for MetroCluster IP configurations

You should verify that your MetroCluster IP configuration and network meets all Inter-Switch Link (ISL) requirements. Although certain requirements might not apply to your configuration, you should still be aware of all of the ISL requirements to gain a better understanding of the overall configuration.

The following table provides an overview of the topics covered in this section.

Title	Description
NetApp-validated and MetroCluster-compliant switches	Describes the switch requirements. Applies to all switches used in MetroCluster configurations, including backend switches.
Considerations for ISLs	Describes the ISL requirements. Applies to all MetroCluster configurations, regardless of network topology and whether you use NetApp-validated switches or MetroCluster-compliant switches.
Considerations when deploying MetroCluster in a shared layer 2 or layer 3 networks	Describes the requirements for shared layer 2 or layer 3 networks. Applies to all configurations except for MetroCluster configurations using NetApp-validated switches and using direct connected ISLs.
Considerations when using MetroCluster Compliant switches	Describes the requirements for MetroCluster-compliant switches. Applies to all MetroCluster configurations that are not using NetApp-validated switches.
Examples of MetroCluster network topologies	Provides examples of different MetroCluster network topologies. Applies to all MetroCluster configurations.

NetApp-validated and MetroCluster-compliant switches in a MetroCluster IP configuration

All of the switches used in your configuration, including backend switches, must either be NetApp-validated or MetroCluster-compliant.

NetApp-validated switches

A switch is NetApp-validated if it meets the following requirements:

- The switch is provided by NetApp as part of the MetroCluster IP configuration
- The switch is listed in the [NetApp Hardware Universe](#) as a supported switch under *MetroCluster-over-IP-connections*
- The switch is only used to connect MetroCluster IP controllers and, in some configurations, NS224 drive shelves
- The switch is configured using the Reference Configuration File (RCF) provided by NetApp

Any switch that does not meet these requirements is **not** a NetApp-validated switch.

MetroCluster-compliant switches

A MetroCluster-compliant switch is not NetApp-validated but can be used in a MetroCluster IP configuration if it meets certain requirements and configuration guidelines.



NetApp does not provide troubleshooting or configuration support services for any non-validated MetroCluster-compliant switch.

Requirements for Inter-Switch Links (ISLs) on MetroCluster IP configurations

Inter-Switch Links (ISLs) carrying MetroCluster traffic on all MetroCluster IP configurations and network topologies have certain requirements. These requirements apply to all ISLs carrying MetroCluster traffic, regardless of whether the ISLs are direct or shared between customer switches.

MetroCluster ISL requirements

The following applies to ISLs on all MetroCluster IP configurations:

- Both fabrics must have the same number of ISLs.
- ISLs on one fabric must all be the same speed and length.
- ISLs in both fabrics must be the same speed and length.
- The maximum supported difference in distance between fabric 1 and fabric 2 is 20km or 0.2ms.
- The ISLs must have the same topology. For example, they should all be direct links, or if the configuration uses WDM, then they must all use WDM.
- The minimum required ISL speed depends on the platform model:
 - Beginning with ONTAP 9.18.1, platforms with a MetroCluster IP backend port speed of 100G require a minimum ISL link speed of 100Gbps. Using a different ISL speed requires a Feature Variance Request (FPVR). To file an FPVR, please contact your NetApp sales team.
 - On all other platforms, the minimum supported ISL link speed is 10Gbps.

- There must be at least one 10Gbps ISL port per fabric.

Latency and packet loss limits in the ISLs

The following applies to round-trip traffic between the MetroCluster IP switches at site_A and site_B, with the MetroCluster configuration in steady state operation:

- As the distance between two MetroCluster sites increases, latency increases, usually in the range of 1 ms round-trip delay time per 100 km (62 miles). Latency also depends on the network service level agreement (SLA) in terms of the bandwidth of the ISL links, packet drop rate, and jitter on the network. Low bandwidth, high jitter, and random packet drops lead to different recovery mechanisms by the switches, or the TCP engine on the controller modules, for successful packet delivery. These recovery mechanisms can increase overall latency. For specific information on round trip latency and maximum distance requirements for your configuration, refer to the [Hardware Universe](#).
- Any device that contributes to latency must be accounted for.
- The [Hardware Universe](#) provides the distance in km. You must allocate 1ms for every 100km. The maximum distance is defined by what is reached first, either the maximum round-trip time (RTT) in ms, or the distance in km. For example – if *The Hardware Universe* lists a distance of 300km, translating to 3ms, your ISL can be no further than 300km and the max RTT cannot exceed 3ms – whichever is reached first.
- Packet loss must be less than, or equal to, 0.01%. The maximum packet loss is the sum of all loss on all links on the path between the MetroCluster nodes, and the loss on the local MetroCluster IP interfaces.
- The supported jitter value is 3ms for round trip (or 1.5ms for one-way).
- The network should allocate and maintain the SLA amount of bandwidth required for MetroCluster traffic, regardless of microbursts and spikes in the traffic.
- If you are using ONTAP 9.7 or later, the intermediate network between the two sites must provide a minimum bandwidth of 4.5Gbps for the MetroCluster IP configuration.

Transceiver and cable considerations

Any SFPs or QSFPs supported by the equipment vendor are supported for the MetroCluster ISLs. SFPs and QSFPs provided by NetApp or the equipment vendor must be supported by the switch and switch firmware.

When connecting the controllers to the switches and the local cluster ISLs, you must use the transceivers and cables provided by NetApp with the MetroCluster.

When you use a QSFP-SFP adapter, whether you configure the port in breakout or native speed mode depends on the switch model and firmware. For example, using a QSFP-SFP adapter with Cisco 9336C switches running NX-OS firmware 9.x or 10.x requires that you configure the port in native speed mode.



If you configure an RCF, verify that you select the correct speed mode or use a port with an appropriate speed mode.

Using xWDM, TDM, and external encryption devices

When you use xWDM/TDM devices or devices providing encryption in a MetroCluster IP configuration your environment must meet the following requirements:

- When connecting the MetroCluster IP switches to the xWDM/TDM, the external encryption devices or xWDM/TDM equipment must be certified by the vendor for the switch and firmware. The certification must cover the operating mode (such as trunking and encryption).
- The overall end-to-end latency and jitter, including the encryption, cannot be more than the maximum

amount stated in the IMT and in this documentation.

Supported number of ISLs and breakout cables

The following table shows the supported maximum number of ISLs that can be configured on a MetroCluster IP switch using the Reference Configuration File (RCF) configuration.

MetroCluster IP switch model	Port type	Maximum number of ISLs
Broadcom-supported BES-53248 switches	Native ports	4 ISLs using 10Gbps or 25Gbps
Broadcom-supported BES-53248 switches	Native ports (Note 1)	2 ISLs using 40Gbps or 100Gbps
Cisco 3132Q-V	Native ports	6 ISLs using 40Gbps
Cisco 3132Q-V	Breakout cables	16 ISLs using 10Gbps
Cisco 3232C	Native ports	6 ISLs using 40Gbps or 100Gbps
Cisco 3232C	Breakout cables	16 ISLs using 10Gbps or 25Gbps
Cisco 9336C-FX2 (not connecting NS224 shelves)	Native ports	6 ISLs using 40Gbps or 100Gbps
Cisco 9336C-FX2 (not connecting NS224 shelves)	Breakout cables	16 ISLs using 10Gbps or 25Gbps
Cisco 9336C-FX2 (connecting NS224 shelves)	Native ports (Note 2)	4 ISLs using 40Gbps or 100Gbps
Cisco 9336C-FX2 (connecting NS224 shelves)	Breakout cables (Note 2)	16 ISLs using 10Gbps or 25Gbps
NVIDIA SN2100	Native ports (Note 2)	2 ISLs using 40Gbps or 100Gbps
NVIDIA SN2100	Breakout cables (Note 2)	8 ISLs using 10Gbps or 25Gbps

Note 1: Using 40Gbps or 100Gbps ISLs on a BES-53248 switch requires an additional license.

Note 2: The same ports are used for native speed and breakout mode. You must choose to use ports in native speed mode or breakout mode when creating the RCF file.

- All ISLs on one MetroCluster IP switch must be the same speed. Using a mix of ISL ports with different speeds concurrently is not supported.
- For optimum performance, you should use at least one 40Gbps ISL per network. You should not use a single 10Gbps ISL per network for FAS9000, AFF A700, or other high capacity platforms.



NetApp recommends that you configure a small number of high bandwidth ISLs, rather than a high number of low bandwidth ISLs. For example, configuring one 40Gbps ISL instead of four 10Gbps ISLs is preferred. When using multiple ISLs, statistical load-balancing can impact the maximum throughput. Uneven balancing can reduce throughput to that of a single ISL.

Requirements to deploy MetroCluster IP configurations in shared layer 2 or layer 3 networks

Depending on your requirements, you can use shared layer 2 or layer 3 networks to deploy MetroCluster.

Beginning with ONTAP 9.6, MetroCluster IP configurations with supported switches can share existing networks for Inter-Switch Links (ISLs) instead of using dedicated MetroCluster ISLs. This topology is known as *shared layer 2 networks*.

Beginning with ONTAP 9.9.1, MetroCluster IP configurations can be implemented with IP-routed (layer 3) backend connections. This topology is known as *shared layer 3 networks*.



- Not all features are supported in all network topologies.
- You must verify that you have adequate network capacity and that the ISL size is appropriate for your configuration. Low latency is critical for replication of data between the MetroCluster sites. Latency issues on these connections can impact client I/O.
- All references to MetroCluster backend switches refer to switches that are NetApp-validated switches or MetroCluster-compliant. See [NetApp-validated and MetroCluster-compliant switches](#) for more details.

ISL requirements for layer 2 and layer 3 networks

The following applies to layer 2 and layer 3 networks:

- The speed and number of ISLs between the MetroCluster switches and the intermediate network switches does not need to match. Similarly, the speed between the intermediate network switches does not need to match.

For example, MetroCluster switches can connect using one 40Gbps ISL to the intermediate switches, and the intermediate switches can connect to each other using two 100Gbps ISLs.

- Network monitoring should be configured on the intermediate network to monitor the ISLs for utilization, errors (drops, link flaps, corruption, and so on), and failures.
- The MTU size must be set to 9216 on all ports carrying MetroCluster end-to-end traffic.
- No other traffic can be configured with a higher priority than class of service (COS) 5.
- Explicit congestion notification (ECN) must be configured on all paths carrying end-to-end MetroCluster traffic.
- ISLs carrying MetroCluster traffic must be native links between the switches.

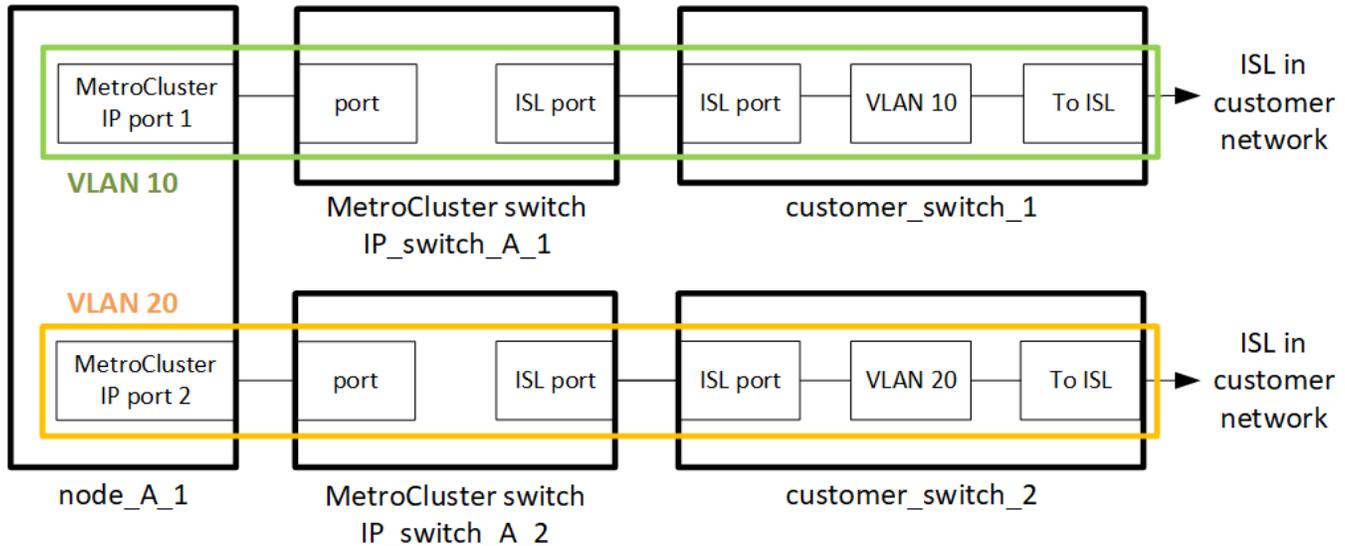
Link sharing services such as Multiprotocol Label Switching (MPLS) links are not supported.

- The layer 2 VLANs must natively span the sites. VLAN overlay such as Virtual Extensible LAN (VXLAN) is not supported.
- The number of intermediate switches is not limited. However, NetApp recommends that you keep the number of switches to the minimum required.

- ISLs on MetroCluster switches are configured with the following:
 - Switch port mode 'trunk' as part of an LACP port-channel
 - The MTU size is 9216
 - No native VLAN is configured
 - Only VLANs carrying cross site MetroCluster traffic are allowed
 - The switch default VLAN is not allowed

Considerations for layer 2 networks

The MetroCluster backend switches are connected to the customer network.



The intermediate customer-provided switches must meet the following requirements:

- The intermediate network must provide the same VLANs between the sites. This must match the MetroCluster VLANs set in the RCF file.
- The RcfFileGenerator does not allow the creation of an RCF file using VLANs that are not supported by the platform.
- The RcfFileGenerator might restrict the use of certain VLAN IDs, for example, if they are intended for future use. Generally, reserved VLANs are up to and including 100.
- Layer 2 VLANs with IDs that match the MetroCluster VLAN IDs must span the shared network.

VLAN configuration in ONTAP

You can only specify the VLAN during interface creation. You can configure the default VLANs 10 and 20, or VLANs within the range 101 to 4096 (or the number supported by the switch vendor, whichever is the lower number). After the MetroCluster interfaces are created, you cannot change the VLAN ID.



Some switch vendors might reserve the use of certain VLANs.

The following systems do not require VLAN configuration within ONTAP. The VLAN is specified by the switch port configuration:

- FAS8200 and AFF A300

- AFF A320
- FAS9000 and AFF A700
- AFF A800, ASA A800, AFF C800, and ASA C800



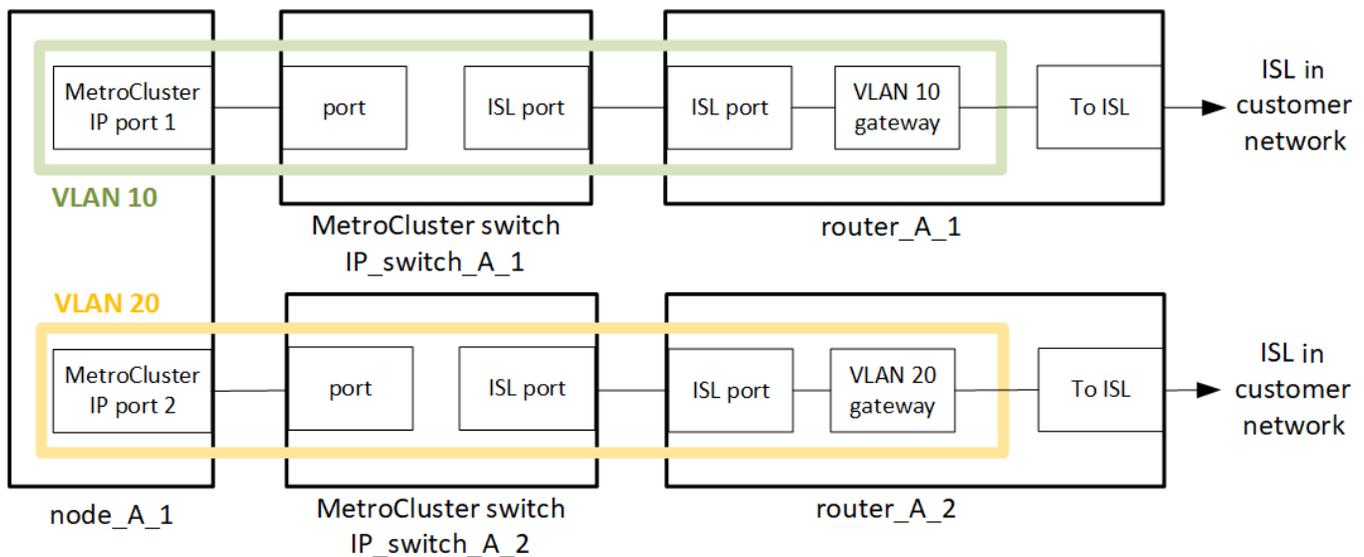
The systems listed above might be configured using VLANs 100 and below. However, some VLANs in this range might be reserved for other or future use.

For all other systems, you must configure the VLAN when you create the MetroCluster interfaces in ONTAP. The following restrictions apply:

- The default VLAN is 10 and 20
- If you are running ONTAP 9.7 or earlier, you can only use the default VLAN 10 and 20.
- If you are running ONTAP 9.8 or later, you can use the default VLAN 10 and 20, and a VLAN over 100 (101 and higher) can also be used.

Considerations for layer 3 networks

The MetroCluster backend switches are connected to the routed IP network, either directly to routers (as shown in the following simplified example) or through other intervening switches.



The MetroCluster environment is configured and cabled as a standard MetroCluster IP configuration as described in [Configure the MetroCluster hardware components](#). When you perform the installation and cabling procedure, you must perform the steps specific to a layer 3 configuration. The following applies to layer 3 configurations:

- You can connect MetroCluster switches directly to the router or to one or more intervening switches.
- You can connect MetroCluster IP interfaces directly to the router or to one of the intervening switches.
- The VLAN must be extended to the gateway device.
- You use the `-gateway` parameter to configure the MetroCluster IP interface address with an IP gateway address.
- The VLAN IDs for the MetroCluster VLANs must be the same at each site. However, the subnets can be different.

- Dynamic routing is not supported for the MetroCluster traffic.
- The following features are not supported:
 - Eight-node MetroCluster configurations
 - Refreshing a four-node MetroCluster configuration
 - Transition from MetroCluster FC to MetroCluster IP
- Two subnets are required on each MetroCluster site—one in each network.
- Auto-IP assignment is not supported.

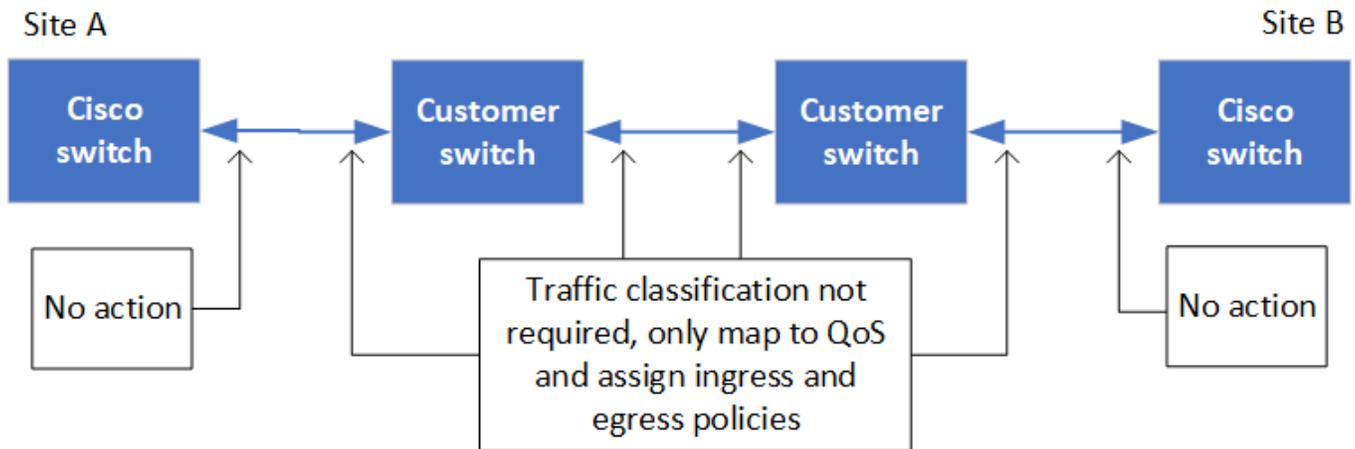
When you configure routers and gateway IP addresses, you must meet the following requirements:

- Two interfaces on one node cannot have the same gateway IP address.
- The corresponding interfaces on the HA pairs on each site must have the same gateway IP address.
- The corresponding interfaces on a node and its DR and AUX partners cannot have the same gateway IP address.
- The corresponding interfaces on a node and its DR and AUX partners must have the same VLAN ID.

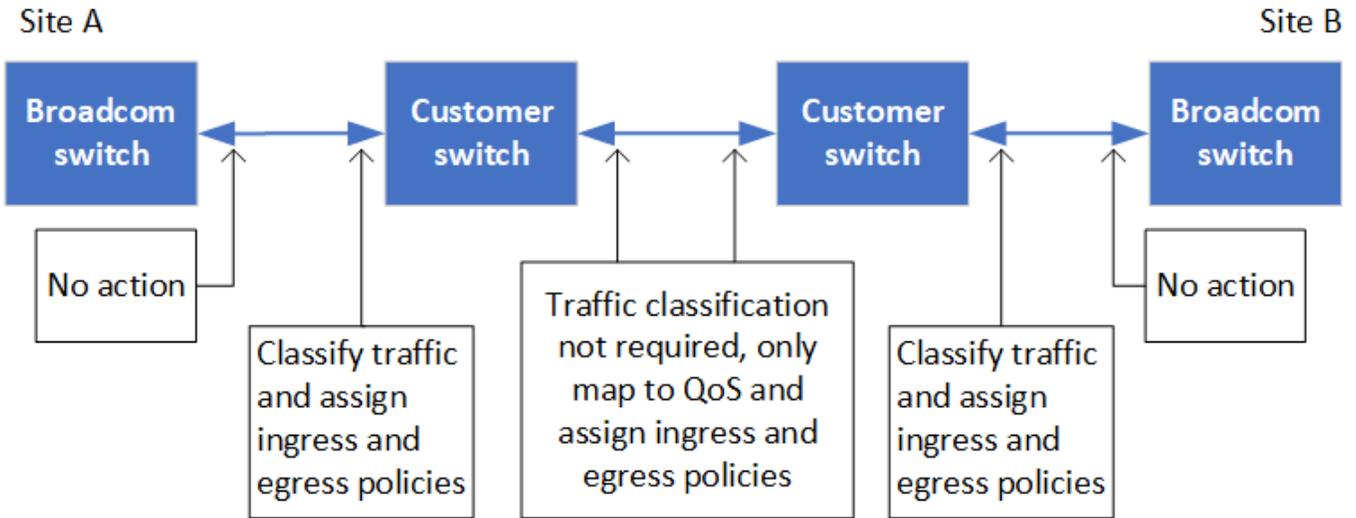
Required settings for intermediate switches

When MetroCluster traffic traverses an ISL in an intermediate network, you should verify that the configuration of the intermediate switches ensures that the MetroCluster traffic (RDMA and storage) meets the required service levels across the entire path between the MetroCluster sites.

The following diagram gives an overview of the required settings when using NetApp validated Cisco switches:



The following diagram gives an overview of the required settings for a shared network when the external switches are Broadcom IP switches.



In this example, the following policies and maps are created for MetroCluster traffic:

- The `MetroClusterIP_ISL_Ingress` policy is applied to ports on the intermediate switch that connects to the MetroCluster IP switches.

The `MetroClusterIP_ISL_Ingress` policy maps the incoming tagged traffic to the appropriate queue on the intermediate switch.

- A `MetroClusterIP_ISL_Egress` policy is applied to ports on the intermediate switch that connect to ISLs between intermediate switches.
- You must configure the intermediate switches with matching QoS access-maps, class-maps, and policy-maps along the path between the MetroCluster IP switches. The intermediate switches map RDMA traffic to COS5 and storage traffic to COS4.

The following examples are for Cisco Nexus 3232C and 9336C-FX2 switches. Depending on your switch vendor and model, you must verify that your intermediate switches have an appropriate configuration.

Configure the class map for the intermediate switch ISL port

The following example shows the class map definitions depending on whether you need to classify or match traffic on ingress.

Classify traffic on ingress:

```
ip access-list rdma
  10 permit tcp any eq 10006 any
  20 permit tcp any any eq 10006
ip access-list storage
  10 permit tcp any eq 65200 any
  20 permit tcp any any eq 65200

class-map type qos match-all rdma
  match access-group name rdma
class-map type qos match-all storage
  match access-group name storage
```

Match traffic on ingress:

```
class-map type qos match-any c5
  match cos 5
  match dscp 40
class-map type qos match-any c4
  match cos 4
  match dscp 32
```

Create an ingress policy map on the ISL port of the intermediate switch:

The following examples show how to create an ingress policy map depending on whether you need to classify or match traffic on ingress.

Classify the traffic on ingress:

```
policy-map type qos MetroClusterIP_ISL_Ingress_Classify
  class rdma
    set dscp 40
    set cos 5
    set qos-group 5
  class storage
    set dscp 32
    set cos 4
    set qos-group 4
  class class-default
    set qos-group 0
```

Match the traffic on ingress:

```
policy-map type qos MetroClusterIP_ISL_Ingress_Match
  class c5
    set dscp 40
    set cos 5
    set qos-group 5
  class c4
    set dscp 32
    set cos 4
    set qos-group 4
  class class-default
    set qos-group 0
```

Configure the egress queuing policy for the ISL ports

The following example shows how to configure the egress queuing policy:

```

policy-map type queuing MetroClusterIP_ISL_Egress
  class type queuing c-out-8q-q7
    priority level 1
  class type queuing c-out-8q-q6
    priority level 2
  class type queuing c-out-8q-q5
    priority level 3
    random-detect threshold burst-optimized ecn
  class type queuing c-out-8q-q4
    priority level 4
    random-detect threshold burst-optimized ecn
  class type queuing c-out-8q-q3
    priority level 5
  class type queuing c-out-8q-q2
    priority level 6
  class type queuing c-out-8q-q1
    priority level 7
  class type queuing c-out-8q-q-default
    bandwidth remaining percent 100
    random-detect threshold burst-optimized ecn

```

These settings must be applied on all switches and ISLs carrying MetroCluster traffic.

In this example, Q4 and Q5 are configured with `random-detect threshold burst-optimized ecn`. Depending on your configuration, you might need to set the minimum and maximum thresholds, as shown in the following example:

```

class type queuing c-out-8q-q5
  priority level 3
  random-detect minimum-threshold 3000 kbytes maximum-threshold 4000
  kbytes drop-probability 0 weight 0 ecn
class type queuing c-out-8q-q4
  priority level 4
  random-detect minimum-threshold 2000 kbytes maximum-threshold 3000
  kbytes drop-probability 0 weight 0 ecn

```



Minimum and maximum values vary depending on the switch and your requirements.

Example 1: Cisco

If your configuration has Cisco switches, you do not need to classify on the first ingress port of the intermediate switch. You then configure the following maps and policies:

- `class-map type qos match-any c5`
- `class-map type qos match-any c4`

- `MetroClusterIP_ISL_Ingress_Match`

You assign the `MetroClusterIP_ISL_Ingress_Match` policy map to the ISL ports carrying MetroCluster traffic.

Example 2: Broadcom

If your configuration has Broadcom switches, you must classify on the first ingress port of the intermediate switch. You then configure the following maps and policies:

- `ip access-list rdma`
- `ip access-list storage`
- `class-map type qos match-all rdma`
- `class-map type qos match-all storage`
- `MetroClusterIP_ISL_Ingress_Classify`
- `MetroClusterIP_ISL_Ingress_Match`

You assign the `MetroClusterIP_ISL_Ingress_Classify` policy map to the ISL ports on the intermediate switch connecting the Broadcom switch.

You assign the `MetroClusterIP_ISL_Ingress_Match` policy map to the ISL ports on the intermediate switch that is carrying MetroCluster traffic but does not connect the Broadcom switch.

MetroCluster IP configuration network topology examples

Beginning with ONTAP 9.6, some additional network configurations are supported for MetroCluster IP configurations. This section provides some examples of the supported network configurations. Not all of the supported topologies are listed.

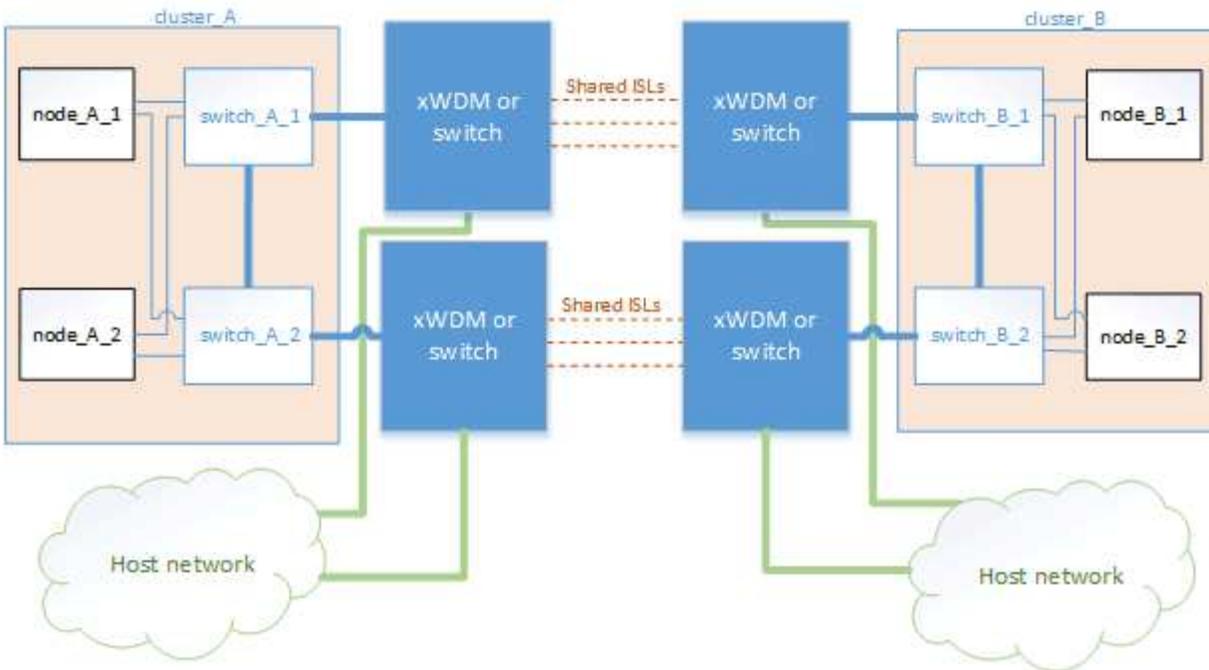
In these topologies, it is assumed that the ISL and intermediate network is configured according to the requirements outlined in [Considerations for ISLs](#).



If you are sharing an ISL with non-MetroCluster traffic, you must verify that the MetroCluster has at least the minimum required bandwidth available at all times.

Shared network configuration with direct links

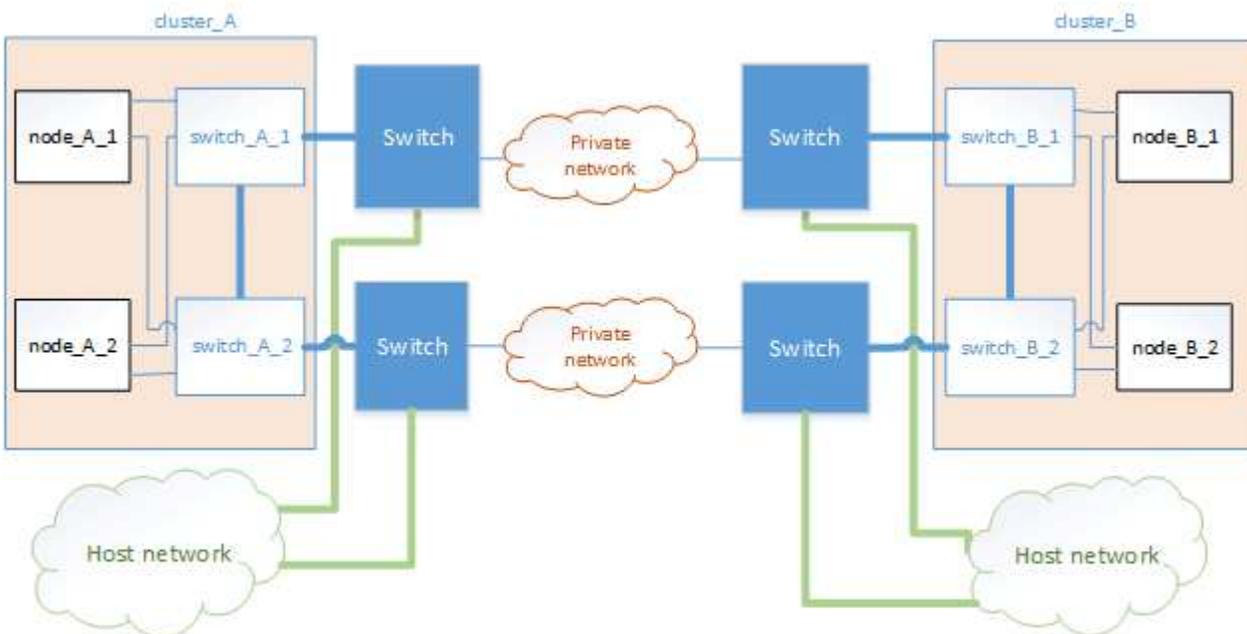
In this topology, two distinct sites are connected by direct links. These links can be between xWDM and TDM devices or switches. The capacity of the ISLs is not dedicated to the MetroCluster traffic but is shared with other non-MetroCluster traffic.



Shared infrastructure with intermediate networks

In this topology, the MetroCluster sites are not directly connected but MetroCluster and the host traffic travel through a network.

The network can consist of a series of xWDM and TDM and switches, but unlike the shared configuration with direct ISLs, the links are not direct between the sites. Depending on the infrastructure between the sites, any combination of network configurations is possible.

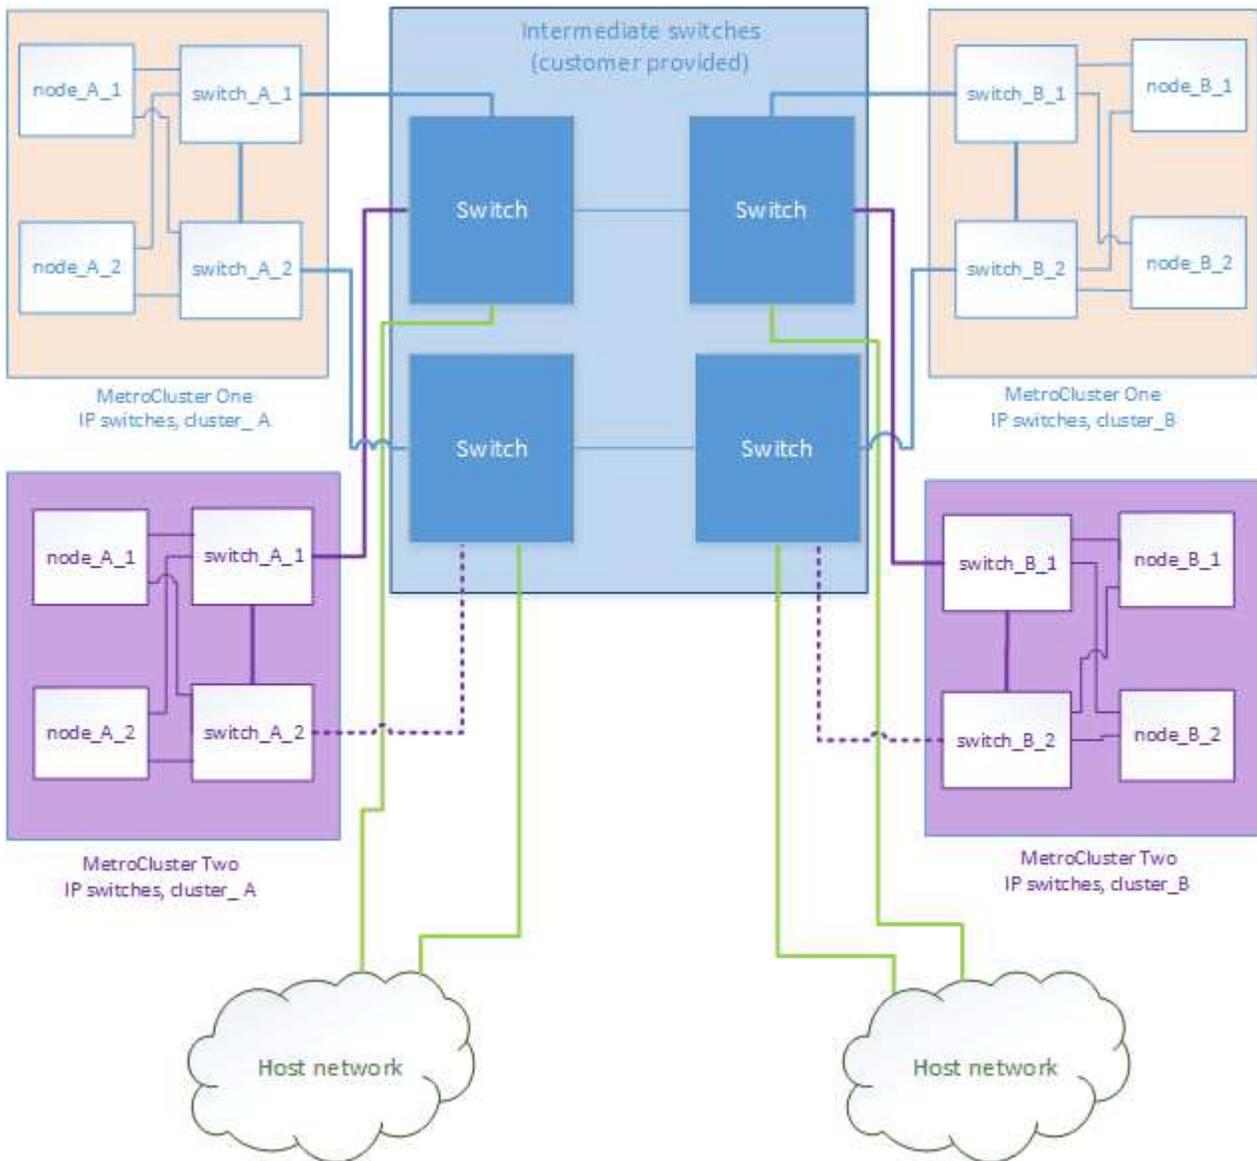


Multiple MetroCluster configurations sharing an intermediate network

In this topology, two separate MetroCluster configurations are sharing the same intermediate network. In the example, MetroCluster one switch_A_1 and MetroCluster two switch_A_1, both connect to the same intermediate switch.

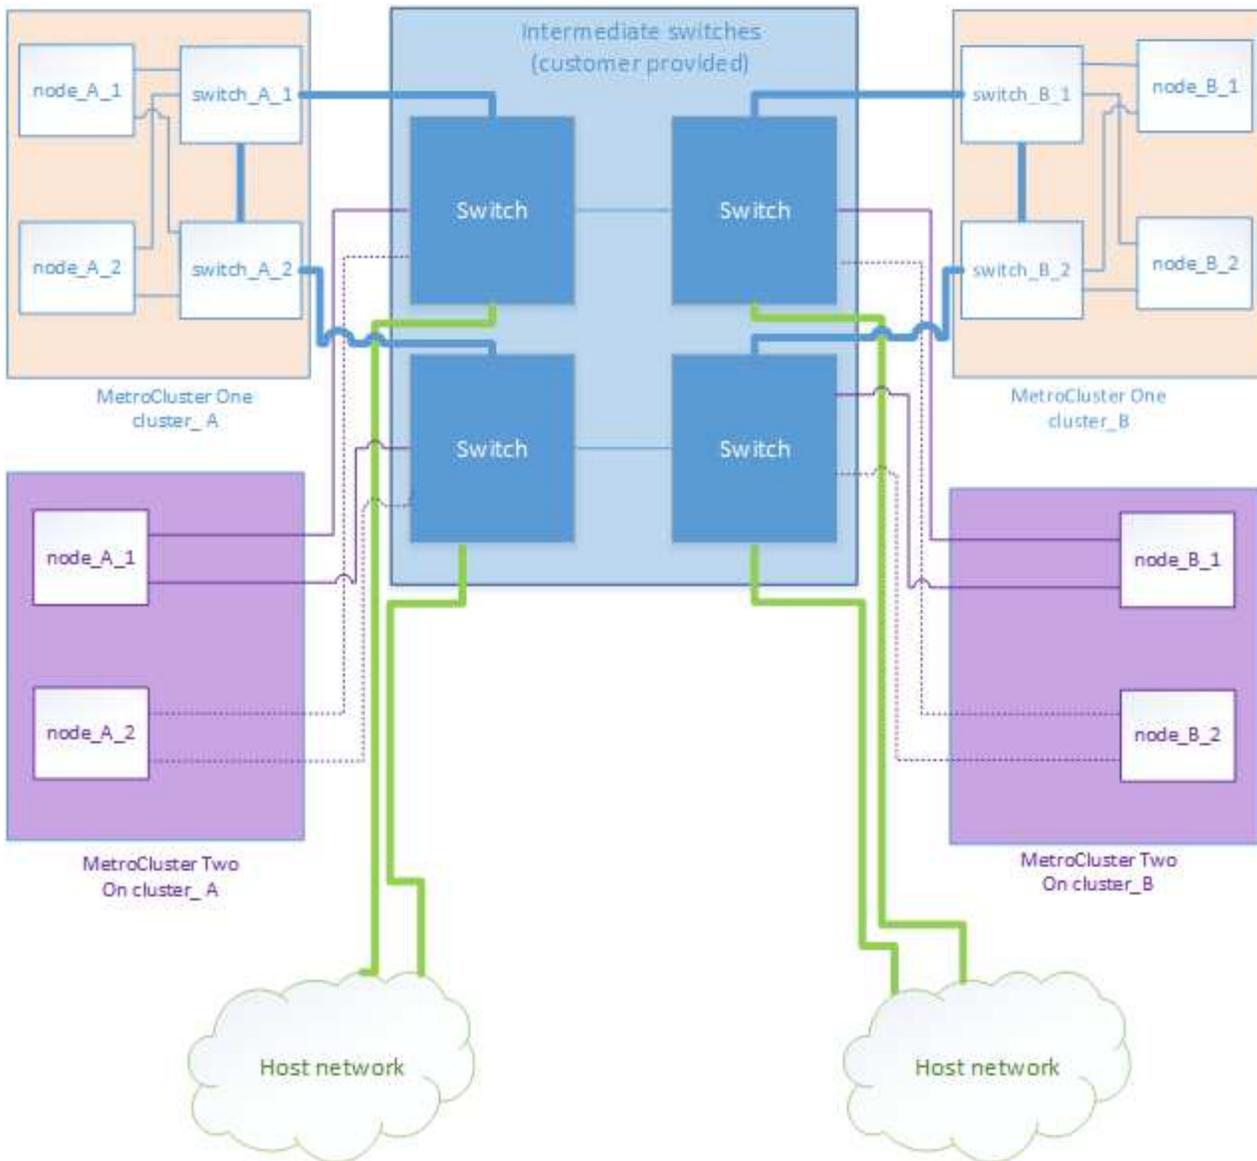


Both “MetroCluster one” or “MetroCluster two” can be one eight-node MetroCluster configuration or two four-node MetroCluster configurations.



Combination of a MetroCluster configuration using NetApp validated switches and a configuration using MetroCluster-compliant switches

Two separate MetroCluster configurations share the same intermediate switch, where one MetroCluster is configured using NetApp validated switches in a shared layer 2 configuration (MetroCluster one), and the other MetroCluster is configured using MetroCluster-compliant switches connecting directly to the intermediate switches (MetroCluster two).



Considerations for using MetroCluster-compliant switches

Requirements and limitations for MetroCluster-compliant switches

Beginning with ONTAP 9.7, MetroCluster IP configurations can use MetroCluster-compliant switches. These are switches that are not NetApp-validated but are compliant with NetApp specifications. However, NetApp does not provide troubleshooting or configuration support services for any non-validated switch. You should be aware of the general requirements and limitations when using MetroCluster-compliant switches.

MetroCluster-compliant versus NetApp-validated switches

A switch is NetApp-validated if it meets the following requirements:

- The switch is provided by NetApp as part of the MetroCluster IP configuration
- The switch is listed in the [NetApp Hardware Universe](#) as a supported switch under *MetroCluster-over-IP-connections*

- The switch is only used to connect MetroCluster IP controllers and, in some configurations, NS224 drive shelves
- The switch is configured using the Reference Configuration File (RCF) provided by NetApp

Any switch that does not meet these requirements is **not** a NetApp-validated switch.

A MetroCluster-compliant switch is not NetApp-validated but can be used in a MetroCluster IP configuration if it meets certain requirements and configuration guidelines.



NetApp does not provide troubleshooting or configuration support services for any non-validated MetroCluster-compliant switch.

General requirements for MetroCluster-compliant switches

The switch connecting the MetroCluster IP interfaces must meet the following general requirements:

- The switches must support quality of service (QoS) and traffic classification.
- The switches must support explicit congestion notification (ECN).
- The switches must support a load-balancing policy to preserve order along the path.
- The switches must support L2 Flow Control (L2FC).
- The switch port must provide a dedicated rate and must not be overallocated.
- The cables and transceivers connecting the nodes to the switches must be provided by NetApp. These cables must be supported by the switch vendor. If you are using optical cabling, the transceiver in the switch might not be provided by NetApp. You must verify that it is compatible with the transceiver in the controller.
- The switches connecting the MetroCluster nodes can carry non-MetroCluster traffic.
- Only platforms that provide dedicated ports for switchless cluster interconnects can be used with a MetroCluster-compliant switch. Platforms such as the FAS2750 and AFF A220 cannot be used because MetroCluster traffic and MetroCluster interconnect traffic share the same network ports.
- The MetroCluster-compliant switch must not be used for local cluster connections.
- The MetroCluster IP interface can be connected to any switch port that can be configured to meet the requirements.
- Four IP switches are required, two for each switch fabric. If you use directors, then you can use a single director at each side, but the MetroCluster IP interfaces must connect to two different blades in two different failure domains on that director.
- The MetroCluster interfaces from one node must connect to two network switches or blades. The MetroCluster interfaces from one node cannot be connected to the same network or switch or blade.
- The network must meet the requirements outlined in the following sections:
 - [Considerations for ISLs](#)
 - [Considerations when deploying MetroCluster in shared layer 2 or layer 3 networks](#)
- The maximum transmission unit (MTU) of 9216 must be configured on all switches that carry MetroCluster IP traffic.
- Reverting to ONTAP 9.6 or earlier is not supported.

Any intermediate switches that you use between the switches connecting the MetroCluster IP interfaces at both sites must meet the requirements and must be configured as outlined in [Considerations when deploying](#)

Limitations when using MetroCluster-compliant switches

You cannot use any configuration or feature that requires that local cluster connections are connected to a switch. For example, you cannot use the following configurations and procedures with a MetroCluster-compliant switch:

- Eight-node MetroCluster configurations
- Transitioning from MetroCluster FC to MetroCluster IP configurations
- Refreshing a four-node MetroCluster IP configuration
- Platforms sharing a physical interface for local cluster and MetroCluster traffic. Refer to [Platform-specific network speeds and switch port modes for MetroCluster-compliant switches](#) for supported speeds.

ONTAP platform-specific network speeds and switch port modes for MetroCluster-compliant switches

If you are using MetroCluster compliant switches, you should be aware of the platform-specific network speeds and switch port mode requirements.

The following table provides platform-specific network speeds and switch port modes for MetroCluster-compliant switches. You should configure the switch port mode according to the table.



- Missing values indicate that the platform cannot be used with a MetroCluster-compliant switch.
- AFF A30, AFF C30, AFF C60, and FAS50 systems require a QSFP-to-SFP+ adapter in the card on the controller to support a 25Gbps network speed.

Platform	Network Speed (Gbps)	Switch port mode
FAS9500 AFF A900 ASA A900	100Gbps 40Gbps when upgrade PCM from FAS9000 / AFF A700	trunk mode
AFF C800 ASA C800 AFF A800 ASA A800	40Gbps or 100Gbps	access mode
FAS9000 AFF A700	40Gbps	access mode
FAS8300 AFF C400 ASA C400 AFF A400 ASA A400	40Gbps or 100Gbps	trunk mode
AFF A320	40Gbps or 100Gbps	access mode
FAS8200 AFF A300	25Gbps	access mode
FAS500f AFF C250 ASA C250 AFF A250 ASA A250	-	-
FAS2750 AFF A220	-	-
AFF A150 ASA A150	-	-
AFF A20	25Gbps	trunk mode
AFF A30	25Gbps or 100Gbps	trunk mode
AFF C30	25Gbps or 100Gbps	trunk mode
AFF C60	25Gbps or 100Gbps	trunk mode
FAS50	25Gbps or 100Gbps	trunk mode
AFF A50	100Gbps	trunk mode
AFF A70	100Gbps	trunk mode
AFF A90	100Gbps	trunk mode
AFF A1K	100Gbps	trunk mode
AFF C80	100Gbps	trunk mode
FAS70	100Gbps	trunk mode
FAS90	100Gbps	trunk mode

MetroCluster IP switch configuration examples

Learn about the various switch port configurations.



The following examples use decimal values and follow the table that applies to Cisco switches. Depending on the switch vendor, you might require different values for DSCP. Refer to the corresponding table for your switch vendor to confirm the correct value.

DSCP value	Decimal	Hex	Meaning
101 000	16	0x10	CS2
011 000	24	0x18	CS3
100 000	32	0x20	CS4
101 000	40	0x28	CS5

Switch port connecting a MetroCluster interface

- Classification for remote direct memory access (RDMA) traffic:
 - Match : TCP port 10006, source, destination, or both
 - Optional match: COS 5
 - Optional match: DSCP 40
 - Set DSCP 40
 - Set COS 5
 - Optional : rate shaping to 20Gbps
- Classification for iSCSI traffic:
 - Match : TCP port 62500, source, destination, or both
 - Optional match: COS 4
 - Optional match: DSCP 32
 - Set DSCP 32
 - Set COS 4
- L2FlowControl (pause), RX and TX

ISL ports

- Classification:
 - Match COS 5 or DSCP 40
 - Set DSCP 40
 - Set COS 5
 - Match COS 4 or DSCP 32
 - Set DSCP 32
 - Set COS 4

- Egress queuing
 - COS group 4 has a minimum configuration threshold of 2000 and a maximum threshold of 3000
 - COS group 5 has a minimum configuration threshold of 3500 and a maximum threshold of 6500.



Configuration thresholds can vary depending on the environment. You must evaluate the configuration thresholds based on your individual environment.

- ECN enabled for Q4 and Q5
- RED enabled for Q4 and Q5

Bandwidth allocation (switch ports connecting MetroCluster interfaces and ISL ports)

- RDMA, COS 5 / DSCP 40: 60%
- iSCSI, COS 4 / DSCP 32: 40%
- Minimum capacity requirement per MetroCluster configuration and network: 10Gbps



If you use rate limits, the traffic should be **shaped** without introducing loss.

Examples for configuring switch ports connecting the MetroCluster controller

The example commands provided are valid for Cisco NX3232 or Cisco NX9336 switches. Commands vary according to the switch type.

If a feature or its equivalent shown in the examples is not available on the switch, the switch does not meet the minimum requirements and cannot be used to deploy a MetroCluster configuration. This is true for any switch connecting to a MetroCluster configuration and for all intermediate switches.



The following examples might only show the configuration for one network.

Basic configuration

A virtual LAN (VLAN) in each network must be configured. The following example shows how to configure a VLAN in network 10.

Example:

```
# vlan 10
The load balancing policy should be set so that order is preserved.
```

Example:

```
# port-channel load-balance src-dst ip-l4port-vlan
```

Examples for configuring classification

You must configure access and class maps to map RDMA and iSCSI traffic to the appropriate classes.

In the following example, all TCP traffic to and from the port 65200 is mapped to the storage (iSCSI) class. All TCP traffic to and from the port 10006 is mapped to the RDMA class. These policy-maps are used on switch

ports connecting the MetroCluster interfaces.

Example:

```
ip access-list storage
 10 permit tcp any eq 65200 any
 20 permit tcp any any eq 65200
ip access-list rdma
 10 permit tcp any eq 10006 any
 20 permit tcp any any eq 10006

class-map type qos match-all storage
 match access-group name storage
class-map type qos match-all rdma
 match access-group name rdma
```

You must configure an ingress policy. An ingress policy maps the traffic as classified to different COS groups. In this example, the RDMA traffic is mapped to COS group 5 and iSCSI traffic is mapped to COS group 4. The ingress policy is used on switch ports connecting the MetroCluster interfaces and on the ISL ports carrying MetroCluster traffic.

Example:

```
policy-map type qos MetroClusterIP_Node_Ingress
class rdma
 set dscp 40
 set cos 5
 set qos-group 5
class storage
 set dscp 32
 set cos 4
 set qos-group 4
```

NetApp recommends that you shape traffic on switch ports connecting a MetroCluster interface, as shown in the following example:

Example:

```
policy-map type queuing MetroClusterIP_Node_Egress
class type queuing c-out-8q-q7
  priority level 1
class type queuing c-out-8q-q6
  priority level 2
class type queuing c-out-8q-q5
  priority level 3
  shape min 0 gbps max 20 gbps
class type queuing c-out-8q-q4
  priority level 4
class type queuing c-out-8q-q3
  priority level 5
class type queuing c-out-8q-q2
  priority level 6
class type queuing c-out-8q-q1
  priority level 7
class type queuing c-out-8q-q-default
  bandwidth remaining percent 100
  random-detect threshold burst-optimized ecn
```

Examples for configuring the node ports

You might need to configure a node port in breakout mode. In the following example, ports 25 and 26 are configured in 4 x 25Gbps breakout mode.

Example:

```
interface breakout module 1 port 25-26 map 25g-4x
```

You might need to configure the MetroCluster interface port speed. The following example shows how to configure the speed to **auto** or into 40Gbps mode:

Example:

```
speed auto

speed 40000
```

The following example shows a switch port configured to connect a MetroCluster interface. It is an access mode port in VLAN 10, with an MTU of 9216 and is operating in native speed. It has symmetric (send and receive) flow control (pause) enabled and the MetroCluster ingress and egress policies assigned.

Example:

```
interface eth1/9
description MetroCluster-IP Node Port
speed auto
switchport access vlan 10
spanning-tree port type edge
spanning-tree bpduguard enable
mtu 9216
flowcontrol receive on
flowcontrol send on
service-policy type qos input MetroClusterIP_Node_Ingress
service-policy type queuing output MetroClusterIP_Node_Egress
no shutdown
```

On 25Gbps ports, you might need to set the Forward Error Correction (FEC) setting to "off", as shown in the following example.

Example:

```
fec off
```

Examples of configuration of ISL ports throughout the network

A MetroCluster-compliant switch is regarded as an intermediate switch, even it directly connects the MetroCluster interfaces. The ISL ports carrying MetroCluster traffic on the MetroCluster-compliant switch must be configured the same way as the ISL ports on an intermediate switch. Refer to [Required settings on intermediate switches](#) for guidance and examples.



Some policy maps are the same for switch ports connecting MetroCluster interfaces and ISLs carrying MetroCluster traffic. You can use the same policy map for both of these port usages.

Learn about unmirrored aggregates in MetroCluster IP configurations

If your configuration includes unmirrored aggregates, you must be aware of potential access issues after switchover operations.

Unmirrored aggregates and hierarchical namespaces

If you are using hierarchical namespaces, you should configure the junction path so that all of the volumes in that path are either on mirrored aggregates only or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates in the junction path might prevent access to the unmirrored aggregates after the switchover operation.

Unmirrored aggregates and maintenance that requires power shutdown

If you perform a negotiated switchover for maintenance that requires a site-wide power shutdown, you should first manually offline any unmirrored aggregates owned by the disaster site.

If you don't offline the unmirrored aggregates owned by the disaster site, nodes at the surviving site might go

down due to multi-disk panics. This might occur if switched-over unmirrored aggregates go offline or are missing because of the loss of connectivity to storage at the disaster site if there's a power shutdown or loss of ISLs.

Unmirrored aggregates, CRS metadata volumes, and data SVM root volumes

The configuration replication service (CRS) metadata volume and data SVM root volumes must be on a mirrored aggregate. You cannot move these volumes to an unmirrored aggregate. If they are on an unmirrored aggregate, negotiated switchover and switchback operations are vetoed and the `metrocluster check` command returns a warning.

Unmirrored aggregates and SVMs

You should configure SVMs on mirrored aggregates only or on unmirrored aggregates only. Configuring SVMs on a mix of both unmirrored and mirrored aggregates can result in a switchover operation that exceeds 120 seconds. This can lead to a data outage if the unmirrored aggregates don't come online.

Unmirrored aggregates and SAN

Before ONTAP 9.9.1, a LUN should not be located on an unmirrored aggregate. Configuring a LUN on an unmirrored aggregate can result in a switchover operation that exceeds 120 seconds and a data outage.

Add storage shelves for unmirrored aggregates

If you add shelves and want to use them for unmirrored aggregates in a MetroCluster IP configuration, you must do the following:

1. Before starting the procedure to add the shelves, issue the following command:

```
metrocluster modify -enable-unmirrored-aggr-deployment true
```

2. Verify that automatic disk assignment is off:

```
disk option show
```

3. Follow the steps of the procedure to add the shelf.
4. Manually assign all disks from new shelf to the node that will own the unmirrored aggregate or aggregates.
5. Create the aggregates:

```
storage aggregate create
```

6. After completing the procedure, issue the following command:

```
metrocluster modify -enable-unmirrored-aggr-deployment false
```

7. Verify that automatic disk assignment is enabled:

```
disk option show
```

Firewall port requirements for MetroCluster IP configurations

If you are using a firewall at a MetroCluster site, you must ensure access for certain

required ports.

Considerations for firewall usage at MetroCluster sites

If you are using a firewall at a MetroCluster site, you must ensure access for required ports.

The following table shows TCP/UDP port usage in an external firewall positioned between two MetroCluster sites.

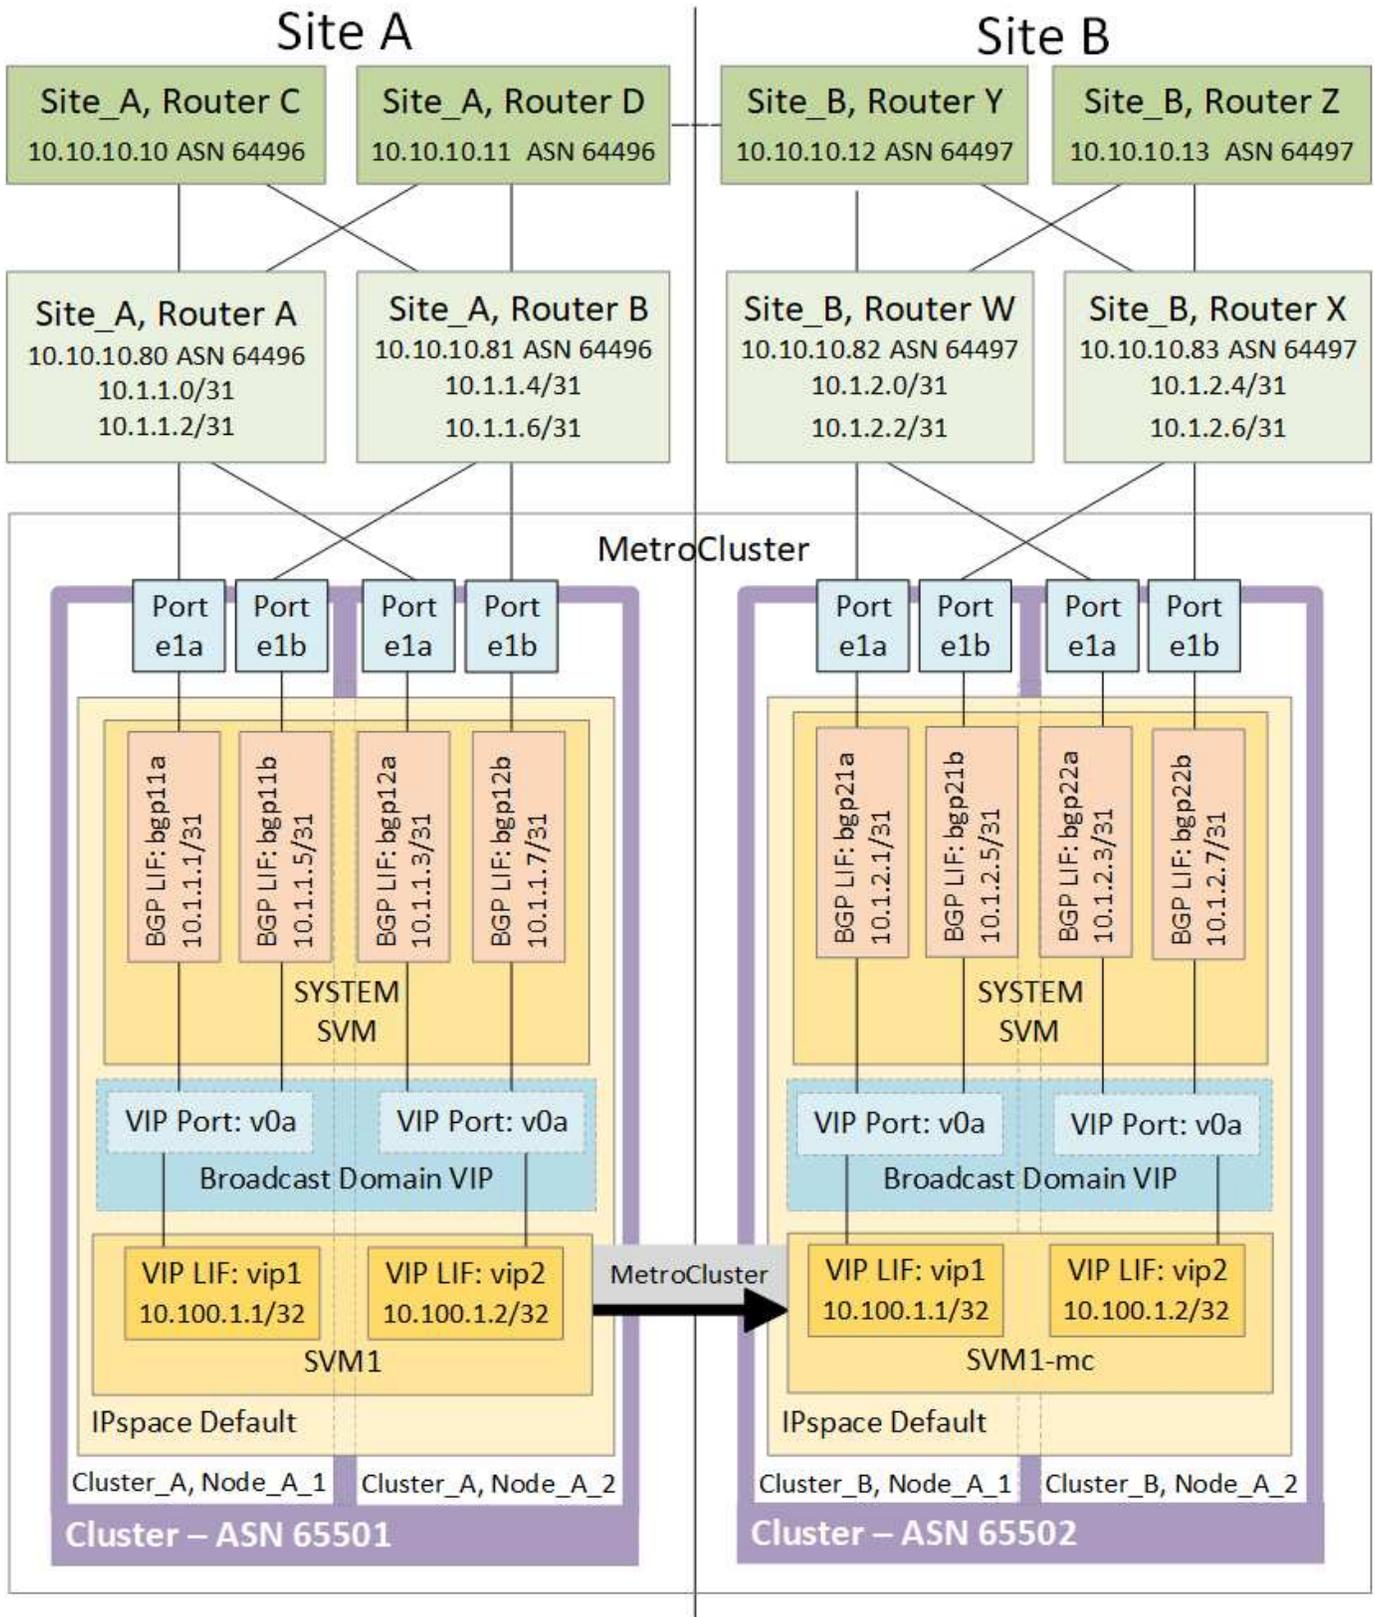
Traffic type	Port/services
Cluster peering	11104 / TCP 11105 / TCP
ONTAP System Manager	443 / TCP
MetroCluster IP intercluster LIFs	65200 / TCP 10006 / TCP and UDP
Hardware assist	4444 / TCP

Learn about using virtual IP and Border Gateway Protocol with a MetroCluster IP configuration

Beginning with ONTAP 9.5, ONTAP supports layer 3 connectivity using virtual IP (VIP) and Border Gateway Protocol (BGP). The combination VIP and BGP for redundancy in the front-end networking with the back-end MetroCluster redundancy provides a layer 3 disaster recovery solution.

Review the following guidelines and illustration when planning your layer 3 solution. For details on implementing VIP and BGP in ONTAP, refer to the following section:

[Configuring virtual IP \(VIP\) LIFs](#)



ONTAP limitations

ONTAP does not automatically verify that all nodes on both sites of the MetroCluster configuration are configured with BGP peering.

ONTAP does not perform route aggregation but announces all individual virtual LIF IPs as unique host routes

at all times.

ONTAP does not support true AnyCast — only a single node in the cluster presents a specific virtual LIF IP (but is accepted by all physical interfaces, regardless of whether they are BGP LIFs, provided the physical port is part of the correct IPspace). Different LIFs can migrate independently of each other to different hosting nodes.

Guidelines for using this Layer 3 solution with a MetroCluster configuration

You must configure your BGP and VIP correctly to provide the required redundancy.

Simpler deployment scenarios are preferred over more complex architectures (for example, a BGP peering router is reachable across an intermediate, non-BGP router). However, ONTAP does not enforce network design or topology restrictions.

VIP LIFs only cover the frontend/data network.

Depending on your version of ONTAP, you must configure BGP peering LIFs in the node SVM, not the system or data SVM. In 9.8, the BGP LIFs are visible in the cluster (system) SVM and the node SVMs are no longer present.

Each data SVM requires the configuration of all potential first hop gateway addresses (typically, the BGP router peering IP address), so that the return data path is available if a LIF migration or MetroCluster failover occurs.

BGP LIFs are node specific, similar to intercluster LIFs — each node has a unique configuration, which does not need to be replicated to DR site nodes.

The existence of the v0a (v0b and so on) continuously validates the connectivity, guaranteeing that a LIF migrate or failover succeeds (unlike L2, where a broken configuration is only visible after the outage).

A major architectural difference is that clients should no longer share the same IP subnet as the VIP of data SVMs. An L3 router with appropriate enterprise grade resiliency and redundancy features enabled (for example, VRRP/HSRP) should be on the path between storage and clients for the VIP to operate correctly.

The reliable update process of BGP allows for smoother LIF migrations because they are marginally faster and have a lower chance of interruption to some clients

You can configure BGP to detect some classes of network or switch misbehaviors faster than LACP, if configured accordingly.

External BGP (EBGP) uses different AS numbers between ONTAP node(s) and peering routers and is the preferred deployment to ease route aggregation and redistribution on the routers. Internal BGP (IBGP) and the use of route reflectors is not impossible but outside the scope of a straightforward VIP setup.

After deployment, you must check that the data SVM is accessible when the associated virtual LIF is migrated between all nodes on each site (including MetroCluster switchover) to verify the correct configuration of the static routes to the same data SVM.

VIP works for most IP-based protocols (NFS, SMB, iSCSI).

Configure the MetroCluster hardware components

Learn about hardware component interconnections in a MetroCluster IP configuration

As you plan your MetroCluster IP configuration, you should understand the hardware components and how they interconnect.

Key hardware elements

A MetroCluster IP configuration includes the following key hardware elements:

- Storage controllers

The storage controllers are configured as two two-node clusters.

- IP network

This back-end IP network provides connectivity for two distinct uses:

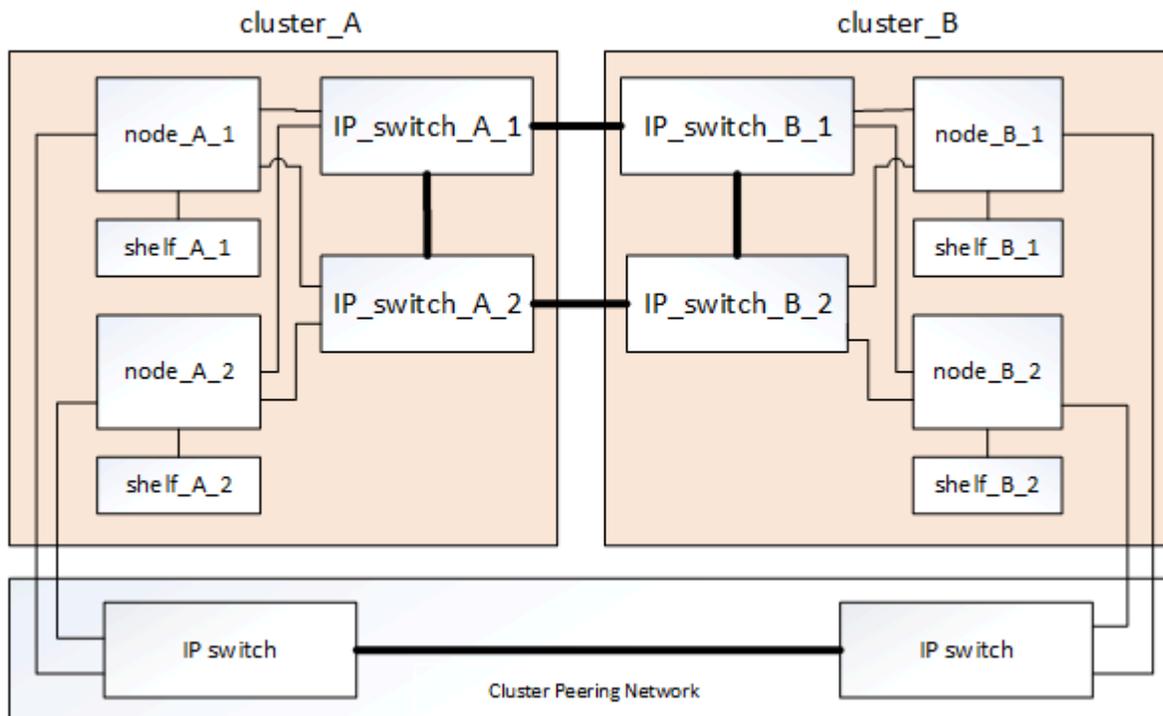
- Standard cluster connectivity for intra-cluster communications.

This is the same cluster switch functionality used in non-MetroCluster switched ONTAP clusters.

- MetroCluster back-end connectivity for replication of storage data and non-volatile cache.

- Cluster peering network

The cluster peering network provides connectivity for mirroring of the cluster configuration, which includes storage virtual machine (SVM) configuration. The configuration of all of the SVMs on one cluster is mirrored to the partner cluster.



Disaster Recovery (DR) groups

A MetroCluster IP configuration consists of one DR group of four nodes.

The following illustration shows the organization of nodes in a four-node MetroCluster configuration:

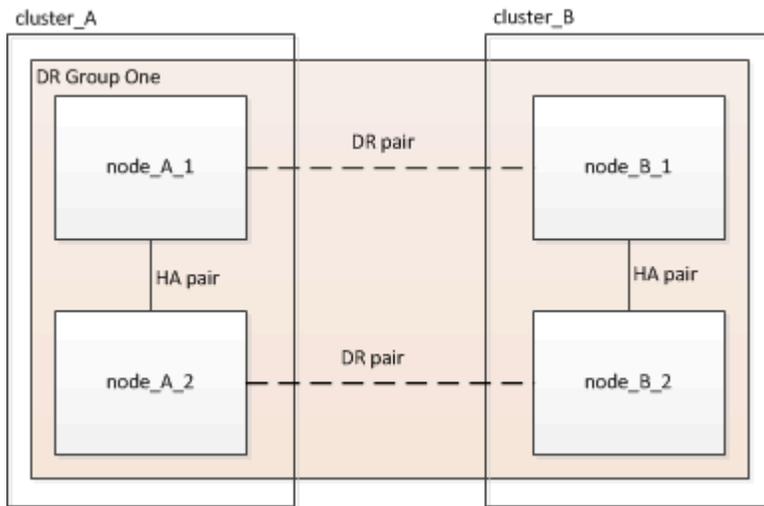
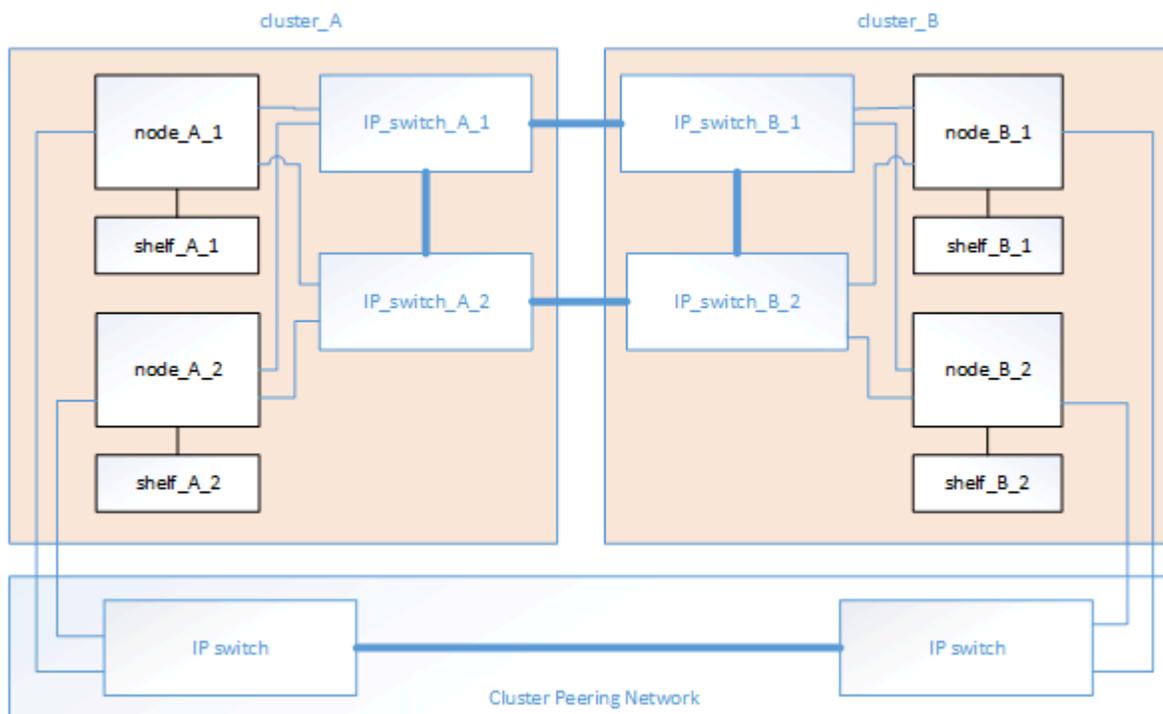


Illustration of the local HA pairs in a MetroCluster configuration

Each MetroCluster site consists of storage controllers configured as an HA pair. This allows local redundancy so that if one storage controller fails, its local HA partner can take over. Such failures can be handled without a MetroCluster switchover operation.

Local HA failover and giveback operations are performed with the storage failover commands, in the same manner as a non-MetroCluster configuration.

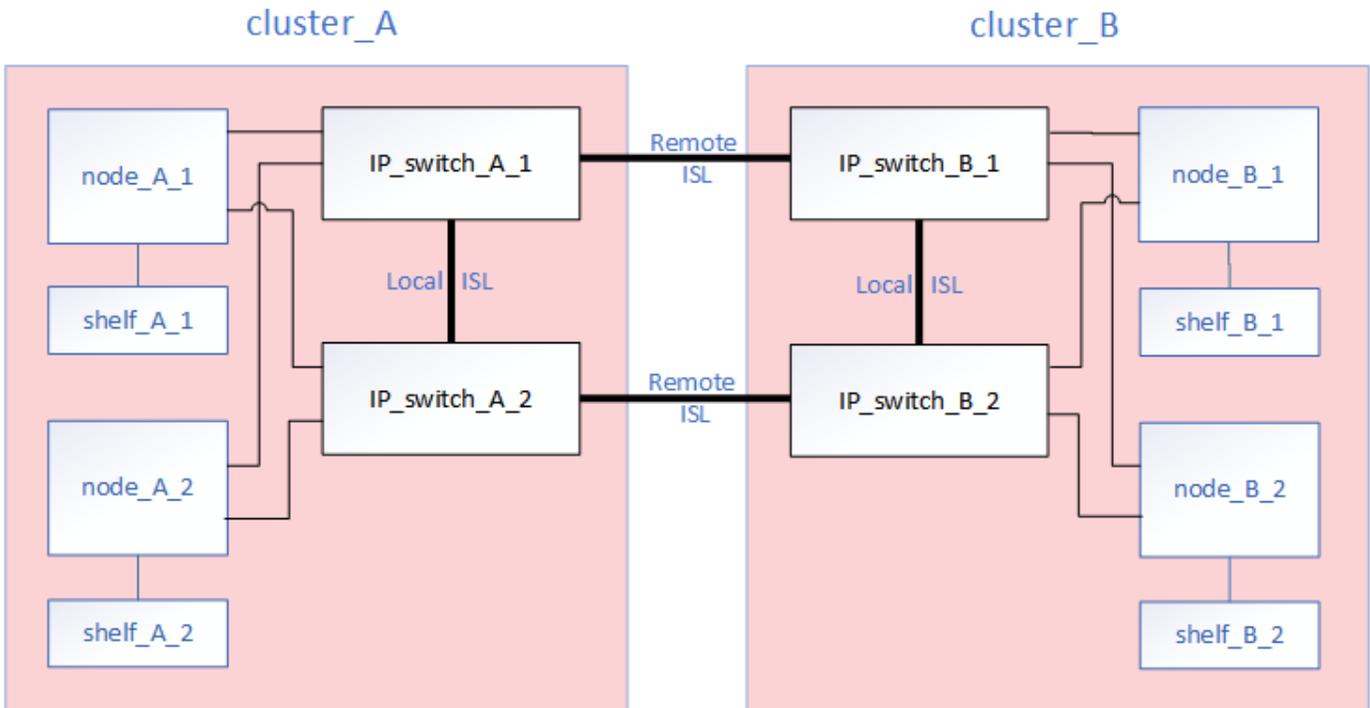


Related information

[ONTAP concepts](#)

Illustration of the MetroCluster IP and cluster interconnect network

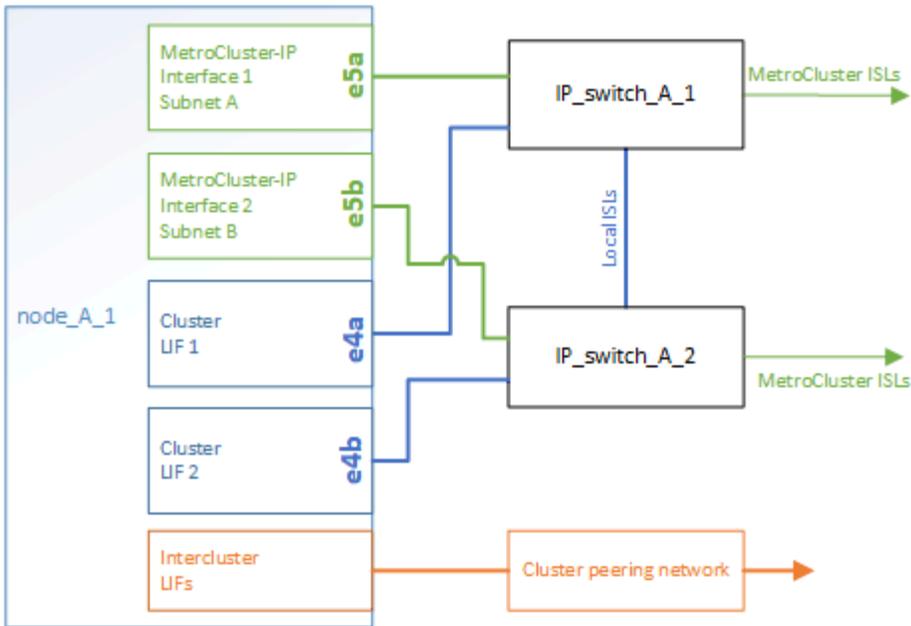
ONTAP clusters typically include a cluster interconnect network for traffic between the nodes in the cluster. In MetroCluster IP configurations, this network is also used for carrying data replication traffic between the MetroCluster sites.



Each node in the MetroCluster IP configuration has dedicated interfaces for connection to the back-end IP network:

- Two MetroCluster IP interfaces
- Two local cluster interfaces

The following illustration shows these interfaces. The port usage shown is for an AFF A700 or FAS9000 system.



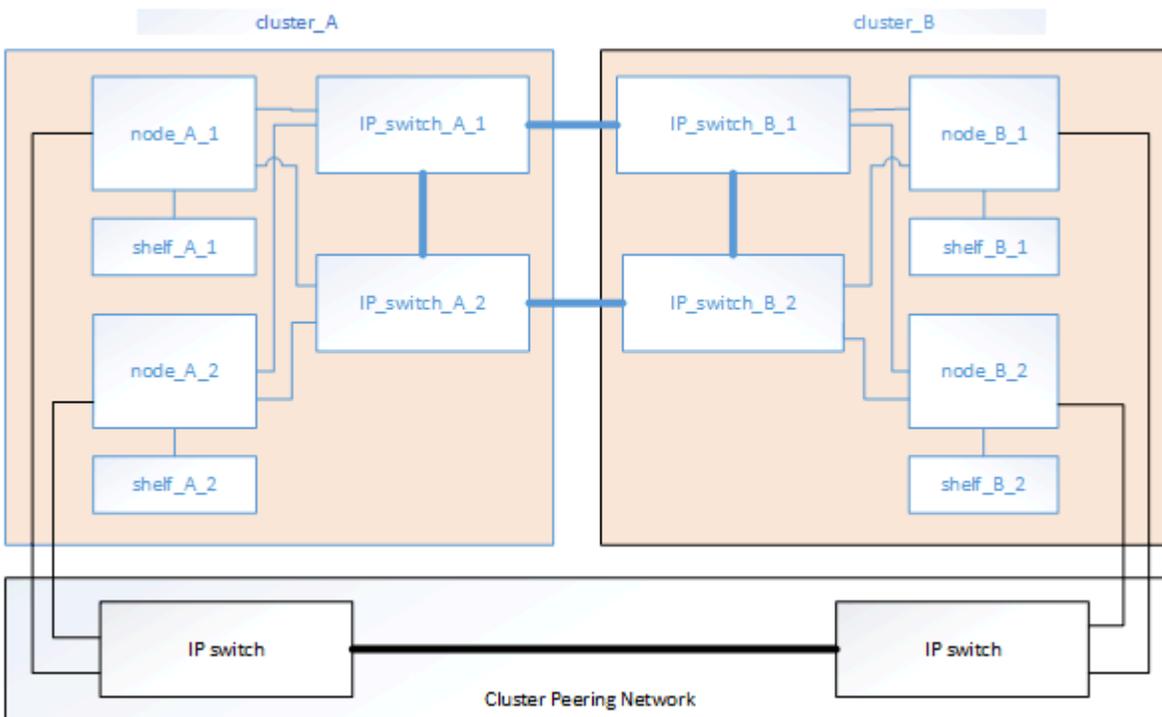
Related information

[Considerations for MetroCluster IP configurations](#)

Illustration of the cluster peering network

The two clusters in the MetroCluster configuration are peered through a customer-provided cluster peering network. Cluster peering supports the synchronous mirroring of storage virtual machines (SVMs, formerly known as Vservers) between the sites.

Intercluster LIFs must be configured on each node in the MetroCluster configuration, and the clusters must be configured for peering. The ports with the intercluster LIFs are connected to the customer-provided cluster peering network. Replication of the SVM configuration is carried out over this network through the Configuration Replication Service.



Related information

[Cluster and SVM peering express configuration](#)

[Considerations for configuring cluster peering](#)

[Cabling the cluster peering connections](#)

[Peering the clusters](#)

Required MetroCluster IP configuration components and naming conventions

Identify the required and supported hardware and software components for a MetroCluster IP configuration. Review the naming conventions that documentation examples use for the components.

Supported software and hardware

The hardware and software must be supported for the MetroCluster IP configuration.

[NetApp Hardware Universe](#)

When using AFF systems, all controller modules in the MetroCluster configuration must be configured as AFF systems.

Hardware redundancy requirements in a MetroCluster IP configuration

Because of the hardware redundancy in the MetroCluster IP configuration, there are two of each component at each site. The sites are arbitrarily assigned the letters A and B, and the individual components are arbitrarily assigned the numbers 1 and 2.

ONTAP cluster requirements in a MetroCluster IP configuration

MetroCluster IP configurations require two ONTAP clusters, one at each MetroCluster site.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A
- Site B: cluster_B

IP switch requirements in a MetroCluster IP configuration

MetroCluster IP configurations require four IP switches. The four switches form two switch storage fabrics that provide the ISL between each of the clusters in the MetroCluster IP configuration.

The IP switches also provide intracluster communication among the controller modules in each cluster.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A

- IP_switch_A_1
- IP_switch_A_2
- Site B: cluster_B
 - IP_switch_B_1
 - IP_switch_B_2

Controller module requirements in a MetroCluster IP configuration

MetroCluster IP configurations require four or eight controller modules.

The controller modules at each site form an HA pair. Each controller module has a DR partner at the other site.

Each controller module must be running the same ONTAP version. Supported platform models depend on the ONTAP version:

- New MetroCluster IP installations on FAS systems are not supported in ONTAP 9.4.
 - Existing MetroCluster IP configurations on FAS systems can be upgraded to ONTAP 9.4.
- Beginning with ONTAP 9.5, new MetroCluster IP installations on FAS systems are supported.
- Beginning with ONTAP 9.4, controller modules configured for ADP are supported.

Example names

The following example names are used in the documentation:

- Site A: cluster_A
 - controller_A_1
 - controller_A_2
- Site B: cluster_B
 - controller_B_1
 - controller_B_2

Gigabit Ethernet adapter requirements in a MetroCluster IP configuration

MetroCluster IP configurations use a 40/100 Gbps or 10/25 Gbps Ethernet adapter for the IP interfaces to the IP switches used for the MetroCluster IP fabric.



Onboard ports are built into the controller hardware (Slot 0) and can't be replaced, so the required slot for adapter is not applicable.

Platform model	Required Gigabit Ethernet adapter	Required slot for adapter	Ports

AFF A900, ASA A900, and FAS9500	X91146A	Slot 5, Slot 7	e5b, e7b Note: Ports e5a and e7a can only be used for intercluster LIFs. You cannot use these ports for a data LIF.
AFF A700 and FAS9000	X91146A-C	Slot 5	e5a, e5b
AFF A800, AFF C800, ASA A800, and ASA C800	X1146A/onboard ports	Slot 1/Not applicable for onboard ports	e0b, e1b
FAS8300, AFF A400, ASA A400, ASA C400, AFF C400	X1146A	Slot 1	e1a, e1b
AFF A300, FAS8200	X1116A	Slot 1	e1a, e1b
FAS2750, AFF A150, ASA A150, AFF A220	Onboard ports	Not applicable	e0a, e0b
FAS500f, AFF A250, ASA A250, ASA C250, AFF C250	Onboard ports	Not applicable	e0c, e0d
AFF A320	Onboard ports	Not applicable	e0g, e0h
AFF A70, FAS70, AFF C80	X50132A	Slot 2	e2a, e2b
AFF A90, AFF A1K, FAS90	X50132A	Slot 2, Slot 3	e2b, e3b Note: Ports e2a and e3a must remain unused. Using these ports for front-end networks or peering is not supported.
AFF A50	X60134A	Slot 2	e2a, e2b
AFF A30, AFF C30, AFF C60, FAS50	X60134A	Slot 2	e2a, e2b
AFF A20	X60132A	Slot 4, Slot 2	e2b, e4b

[Learn about automatic drive assignment and ADP systems in MetroCluster IP configurations.](#)

Pool and drive requirements (minimum supported)

A four-node MetroCluster IP configuration requires the minimum configuration at each site:

- Each node has at least one local pool and one remote pool at the site.
- At least seven drives in each pool.

In a four-node MetroCluster configuration with a single mirrored data aggregate per node, the minimum configuration requires 24 disks at the site.



Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.

In a minimum supported configuration, each pool has the following drive layout:

- Three root drives
- Three data drives
- One spare drive

In a minimum supported configuration, at least one shelf is needed per site.

MetroCluster configurations support RAID-DP, RAID4, and RAID-TEC.



Beginning with ONTAP 9.4, MetroCluster IP configurations support new installations using automatic disk assignment and ADP (Advanced Drive Partitioning). Refer to [Considerations for automatic drive assignment and ADP systems](#) for more information.

Drive location considerations for partially populated shelves

For correct auto-assignment of drives when using shelves that are half populated (12 drives in a 24-drive shelf), drives should be located in slots 0-5 and 18-23.

In a configuration with a partially populated shelf, the drives must be evenly distributed in the four quadrants of the shelf.

Drive location considerations for AFF A800 internal drives

For correct implementation of the ADP feature, the AFF A800 system disk slots must be divided into quarters and the disks must be located symmetrically in the quarters.

An AFF A800 system has 48 drive bays. The bays can be divided into quarters:

- Quarter one:
 - Bays 0 - 5
 - Bays 24 - 29
- Quarter two:
 - Bays 6 - 11
 - Bays 30 - 35
- Quarter three:

- Bays 12 - 17
- Bays 36 - 41
- Quarter four:
 - Bays 18 - 23
 - Bays 42 - 47

If this system is populated with 16 drives, they must be symmetrically distributed among the four quarters:

- Four drives in the first quarter: 0, 1, 2, 3
- Four drives in the second quarter: 6, 7, 8, 9
- Four drives in the third quarter: 12, 13, 14, 15
- Four drives in the fourth quarter: 18, 19, 20, 21

Rack the MetroCluster IP configuration hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

About this task

This task must be performed on both MetroCluster sites.

Steps

1. Plan out the positioning of the MetroCluster components.

The rack space depends on the platform model of the controller modules, the switch types, and the number of disk shelf stacks in your configuration.

2. Properly ground yourself.
3. Install the controller modules in the rack or cabinet.

Follow the steps to *Install Hardware* under the *Install and Setup* instructions for your platform model in the [ONTAP hardware systems documentation](#).

4. Install the IP switches in the rack or cabinet.
5. Install the disk shelves, power them on, and then set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Unique shelf IDs are highly recommended for each SAS disk shelf within each MetroCluster DR group to aid troubleshooting.



Do not cable disk shelves intended to contain unmirrored aggregates at this time. You must wait to deploy shelves intended for unmirrored aggregates until after the MetroCluster configuration is complete and only deploy them after using the `metrocluster modify -enable-unmirrored-aggr-deployment true` command.

Cable the MetroCluster IP switches

How to use the port tables with multiple MetroCluster IP configurations

You must understand how to use the information in the port tables to correctly generate your RCF files.

Before you begin

Review these considerations before using the tables:

- The following tables show the port usage for site A. The same cabling is used for site B.
- You cannot configure the switches with ports of different speeds (for example, a mix of 100 Gbps ports and 40 Gbps ports).
- Keep track of the MetroCluster port group (MetroCluster 1, MetroCluster 2, etc.). You'll need this information when using the RcfFileGenerator tool as described later in this configuration procedure.
- You should cable all of the nodes in the same way. If there are different port combination options available to cable the nodes, all nodes should use the same port combinations. For example, e1a on node1 and e1a on node2 should be attached to one switch. Similarly, the second port from both nodes should be attached to the second switch.
- The [RcfFileGenerator for MetroCluster IP](#) also provides a per-port cabling overview for each switch. Use this cabling overview to verify your cabling.

Cabling two MetroCluster configurations to the switches

When cabling more than one MetroCluster configuration to a switch, you cable each MetroCluster according to the appropriate table. For example, if you are cabling a FAS2750 and an AFF A700 to the same switch, you cable the FAS2750 as per "MetroCluster 1" in Table 1, and the AFF A700 as per "MetroCluster 2" or "MetroCluster 3" in Table 2. You cannot physically cable both the FAS2750 and the AFF A700 as "MetroCluster 1".

Cabling eight-node MetroCluster configurations

For MetroCluster configuration running ONTAP 9.8 and earlier, some procedures that are performed to transition an upgrade require the addition of a second four-node DR group to the configuration to create a temporary eight-node configuration. Beginning with ONTAP 9.9.1, permanent eight-node MetroCluster configurations are supported.

About this task

For eight-node configurations, you use the same method as described above. Instead of a second MetroCluster, you are cabling an additional four-node DR group.

For example, your configuration includes the following:

- Cisco 3132Q-V switches
- MetroCluster 1: FAS2750 platforms
- MetroCluster 2: AFF A700 platforms (these platforms are being added as a second four-node DR group)

Steps

1. For MetroCluster 1, cable the Cisco 3132Q-V switches using the table for the FAS2750 platform and the rows for MetroCluster 1 interfaces.
2. For MetroCluster 2 (the second DR group), cable the Cisco 3132Q-V switches using the table for the AFF A700 platform and the rows for MetroCluster 2 interfaces.

Platform port assignments for Cisco 3132Q-V switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review these guidelines before using the tables:

- If you configure the switch for MetroCluster FC to IP transition, port 5, port 6, port 13, or port 14 can be used to connect the local cluster interfaces of the MetroCluster FC node. Refer to the [RcfFileGenerator](#) and the generated cabling files for more details on cabling this configuration. For all other connections, you can use the port usage assignments listed in the tables.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
FAS2750, AFF A220	Cisco 3132Q-V platform port assignments (group 1)
FAS9000, AFF A700	Cisco 3132Q-V platform port assignments (group 2)
AFF A800, ASA A800	Cisco 3132Q-V platform port assignments (group 3)

Cisco 3132Q-V platform port assignments (group 1)

Review the platform port assignments to cable a FAS2750 or AFF A220 system to a Cisco 3132Q-V switch:

Switch Port	Port use	FAS2750 AFF A220	
		IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled	
7	ISL, Local Cluster native speed / 40G / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b
9/2-4		disabled	
10/1		e0a	e0b
10/2-4		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b
11/2-4		disabled	
12/1		e0a	e0b
12/2-4		disabled	
13/1	MetroCluster 3, Shared Cluster and MetroCluster interface	e0a	e0b
13/2-4		disabled	
14/1		e0a	e0b
14/2-4		disabled	
15	ISL, MetroCluster native speed 40G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25 - 32	Unused	disabled	

Cisco 3132Q-V platform port assignments (group 2)

Review the platform port assignments to cable a FAS9000 or AFF A700 system to a Cisco 3132Q-V switch:

Switch Port	Port use	FAS9000 AFF A700	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a
2			
3	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a
4			
5	MetroCluster 3, Local Cluster interface	e4a	e4e / e8a
6			
7	ISL, Local Cluster native speed 40G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e5a	e5b
10			
11	MetroCluster 2, MetroCluster interface	e5a	e5b
12			
13	MetroCluster 3, MetroCluster interface	e5a	e5b
14			
15	ISL, MetroCluster native speed 40G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25 - 32	Unused	disabled	

Cisco 3132Q-V platform port assignments (group 3)

Review the platform port assignments to cable an AFF A800 or ASA A800 system to a Cisco 3132Q-V switch:

Switch Port	Port use	AFF A800 ASA A800	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e1a
2			
3	MetroCluster 2, Local Cluster interface	e0a	e1a
4			
5	MetroCluster 3, Local Cluster interface	e0a	e1a
6			
7	ISL, Local Cluster native speed 40G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e0b	e1b
10			
11	MetroCluster 2, MetroCluster interface	e0b	e1b
12			
13	MetroCluster 3, MetroCluster interface	e0b	e1b
14			
15	ISL, MetroCluster native speed 40G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25 - 32	Unused	disabled	

Platform port assignments for Cisco 3232C or 36-port Cisco 9336C switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- The tables in this section are for Cisco 3232C switches or 36-port Cisco 9336C-FX2 switches that aren't connecting NS224 storage.

If you have a 12-port Cisco 9336C-FX2 switch, use the tables in [Platform port assignments for 12-port Cisco 9336C-FX2 switches](#).

If you have a 36-port Cisco 9336C-FX2 switch and at least one MetroCluster configuration or DR group is connecting NS224 shelves to the MetroCluster switch, use the tables in [Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage](#).

- The following tables show the port usage for site A. The same cabling is used for site B.

- You cannot configure the switches with ports of different speeds (for example, a mix of 100 Gbps ports and 40 Gbps ports).
- If you are configuring a single MetroCluster with the switches, use the **MetroCluster 1** port group.

Keep track of the MetroCluster port group (MetroCluster 1, MetroCluster 2, MetroCluster 3, or MetroCluster 4). You will need it when using the RcfFileGenerator tool as described later in this configuration procedure.

- The RcfFileGenerator for MetroCluster IP also provides a per-port cabling overview for each switch.

Use this cabling overview to verify your cabling.

- RCF file version v2.10 or later is required for 25G breakout mode for MetroCluster ISLs.
- ONTAP 9.13.1 or later and RCF file version 2.00 are required to use a platform other than FAS8200 or AFF A300 in the "MetroCluster 4" group.



The RCF file version is different to the version of the RCFfilegenerator tool used to generate the file. For example, you can generate an RCF file version 2.00 using RCFfilegenerator v1.6c.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS2750, AFF A220 FAS500f, AFF C250, ASA C250 AFF A250, ASA A250	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 1)
AFF A20	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 2)
AFF A30, AFF C30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 3a) or 100G (group 3b) Ethernet card. <ul style="list-style-type: none"> • Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3a - 25G) • Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3b - 100G)
FAS8200, AFF A300	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 4)
AFF A320 FAS8300, AFF C400, ASA C400, FAS8700 AFF A400, ASA A400	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 5)
AFF A50	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 6)

If your system is...	Use this cabling table...
FAS9000, AFF A700 AFF C800, ASA C800, AFF A800, ASA A800 FAS9500, AFF A900, ASA A900	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 7)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 8)

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS2750, AFF A220, FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A150 ASA A150 FAS2750 AFF A220		FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
9/2-4		disabled		disabled	
10/1		e0a	e0b	e0c	e0d
10/2-4		disabled		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
11/2-4		disabled		disabled	
12/1		e0a	e0b	e0c	e0d
12/2-4		disabled		disabled	
13/1	MetroCluster 3, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
13/2-4		disabled		disabled	
14/1		e0a	e0b	e0c	e0d
14/2-4		disabled		disabled	
15	ISL, MetroCluster native speed 40G / 100G				
16					
17		ISL, MetroCluster		ISL, MetroCluster	
18					
19					
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G				
22/1-4		ISL, MetroCluster		ISL, MetroCluster	
23/1-4					
24/1-4					
25/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
25/2-4		disabled		disabled	
26/1		e0a	e0b	e0c	e0d
26/2-4		disabled		disabled	
27 - 32	Unused	disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 2)

Review the platform port assignments to cable an AFF A20 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e2a	e4a
1/2-4		disabled	
2/1		e2a	e4a
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e2a	e4a
3/2-4		disabled	
4/1		e2a	e4a
4/2-4		disabled	
5/1	MetroCluster 3, Local Cluster interface	e2a	e4a
5/2-4		disabled	
6/1		e2a	e4a
6/2-4		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e2b	e4b
9/2-4		disabled	
10/1		e2b	e4b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e2b	e4b
11/2-4		disabled	
12/1		e2b	e4b
12/2-4		disabled	
13/1	MetroCluster 3, MetroCluster interface	e2b	e4b
13/2-4		disabled	
14/1		e2b	e4b
14/2-4		disabled	
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25/1	MetroCluster 4, MetroCluster interface	e2b	e4b
25/2-4		disabled	
26/1		e2b	e4b
26/2-4		disabled	
27 - 28	Unused	disabled	
29/1	MetroCluster 4, Local Cluster interface	e2a	e4a
29/2-4		disabled	
30/1		e2a	e4a
30/2-4		disabled	
25 - 32	Unused	disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a Cisco 3232C or 9336C-FX2 switch using a four-port 25G Ethernet card.



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled		disabled	
2/1		e4a	e4b	e4a	e4b	e4a	e4b
2/2-4		disabled		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled		disabled	
4/1		e4a	e4b	e4a	e4b	e4a	e4b
4/2-4		disabled		disabled		disabled	
5/1	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
5/2-4		disabled		disabled		disabled	
6/1		e4a	e4b	e4a	e4b	e4a	e4b
6/2-4		disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
14		e2a	e2b	e2a	e2b	e2a	e2b
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
17		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
18		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
19		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
24/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
25	MetroCluster 4, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
26		e2a	e2b	e2a	e2b	e2a	e2b
27 - 28	Unused	disabled		disabled		disabled	
29/1	MetroCluster 4, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
29/2-4		disabled		disabled		disabled	
30/1		e4a	e4b	e4a	e4b	e4a	e4b
30/2-4		disabled		disabled		disabled	
25 - 32	Unused	disabled		disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 3b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a Cisco 3232C or 9336C-FX2 switch using a two-port 100G Ethernet card.



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
		1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
6		e4a	e4b	e4a	e4b	e4a	e4b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
14		e2a	e2b	e2a	e2b	e2a	e2b
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16							
17							
18							
19							
20							
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4							
23/1-4							
24/1-4							
25	MetroCluster 4, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
26		e2a	e2b	e2a	e2b	e2a	e2b
27 - 28	Unused	disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
30		e4a	e4b	e4a	e4b	e4a	e4b
25 - 32	Unused	disabled		disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 4)

Review the platform port assignments to cable a FAS8200 or AFF A300 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	FAS8200 AFF A300	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e0a	e0b
1/2-4		disabled	
2/1		e0a	e0b
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e0a	e0b
3/2-4		disabled	
4/1		e0a	e0b
4/2-4		disabled	
5/1	MetroCluster 3, Local Cluster interface	e0a	e0b
5/2-4		disabled	
6/1		e0a	e0b
6/2-4		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e1a	e1b
9/2-4		disabled	
10/1		e1a	e1b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e1a	e1b
11/2-4		disabled	
12/1		e1a	e1b
12/2-4		disabled	
13/1	MetroCluster 3, MetroCluster interface	e1a	e1b
13/2-4		disabled	
14/1		e1a	e1b
14/2-4		disabled	
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25/1	MetroCluster 4, MetroCluster interface	e1a	e1b
25/2-4		disabled	
26/1		e1a	e1b
26/2-4		disabled	
27 - 28	Unused	disabled	
29/1	MetroCluster 4, Local Cluster interface	e0a	e0b
29/2-4		disabled	
30/1		e0a	e0b
30/2-4		disabled	
25 - 32	Unused	disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled	

If you are upgrading from older RCF files, the cabling configuration might be using ports in the "MetroCluster 4" group (ports 25/26 and 29/30).

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 5)

Review the platform port assignments to cable an AFF A320, FAS8300, AFF C400, ASA C400, FAS8700, AFF A400, or ASA A400 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A320		FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
2							
3	MetroCluster 2, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
4							
5	MetroCluster 3, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
6							
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
10							
11	MetroCluster 2, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
12							
13	MetroCluster 3, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
14							
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16							
17							
18							
19							
20							
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4							
23/1-4							
24/1-4							
25	MetroCluster 4, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
26							
27 - 28	Unused	disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
30							
31 - 32	Unused	disabled		disabled		disabled	
33 - 34	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	



Using ports in the "MetroCluster 4" group requires ONTAP 9.13.1 or later.

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 6)

Review the platform port assignments to cable an AFF A50 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2		e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4		e4a	e4b
5	MetroCluster 3, Local Cluster interface	e4a	e4b
6		e4a	e4b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10		e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12		e2a	e2b
13	MetroCluster 3, MetroCluster interface	e2a	e2b
14		e2a	e2b
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster	
16			
17			
18			
19			
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster	
22/1-4			
23/1-4			
24/1-4			
25	MetroCluster 4, MetroCluster interface	e2a	e2b
26		e2a	e2b
27 - 28	Unused	disabled	
29	MetroCluster 4, Local Cluster interface	e4a	e4b
30		e4a	e4b
25 - 32	Unused	disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled	

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 7)

Review the platform port assignments to cable a FAS9000, AFF A700, AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	FAS9000 AFF A700		AFF C800 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
2							
3							
4	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
5							
6	MetroCluster 3, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
10							
11	MetroCluster 2, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
12							
13	MetroCluster 3, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
14							
15	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
16							
17							
18							
19							
20							
21/1-4	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4							
23/1-4							
24/1-4							
25	MetroCluster 4, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
26							
27 - 28	Unused	disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
30							
31 - 32	Unused	disabled		disabled		disabled	
33 - 34	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled	

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).



Using ports in the "MetroCluster 4" group requires ONTAP 9.13.1 or later.

Cisco 3232C or Cisco 9336C-FX2 platform port assignments (group 8)

Review the platform port assignments to cable an AFF A70, FAS70, AFF C80, FAS90, AFF A90, or AFF A1K system to a Cisco 3232C or 9336C-FX2 switch:

Switch Port	Port use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3									
4	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
5									
6	MetroCluster 3, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
7									
8	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11									
12	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
13									
14	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
15									
16	ISL, MetroCluster native speed 40G / 100G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
17									
18									
19									
20	ISL, MetroCluster breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4									
22/1-4									
23/1-4									
24/1-4									
25	MetroCluster 4, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
26									
27 - 28	Unused	disabled		disabled		disabled		disabled	
29	MetroCluster 4, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
30									
31 - 32	Unused	disabled		disabled		disabled		disabled	
33 - 36	Unused (Cisco 9336C-FX2 only)	disabled		disabled		disabled		disabled	

Platform port assignments for 12-port Cisco 9336C-FX2 switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- The tables in this section are for 12-port Cisco 9336C-FX2 switches.

If you have a 36-port Cisco 9336C-FX2 switch that isn't connecting NS224 shelves, use the tables in [Platform port assignments for Cisco 3232C or 36-port Cisco 9336C-FX2 switches](#).

If you have a 36-port Cisco 9336C-FX2 switch and at least one MetroCluster configuration or DR group is connecting NS224 shelves to the MetroCluster switch, use the tables in [Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage](#).



The 12-port Cisco 9336C-FX2 switch doesn't support connecting NS224 shelves to the MetroCluster switch.

- The following tables show the port usage for site A. The same cabling is used for site B.
- You cannot configure the switches with ports of different speeds (for example, a mix of 100 Gbps ports and 40 Gbps ports).
- If you are configuring a single MetroCluster with the switches, use the **MetroCluster 1** port group.

Keep track of the MetroCluster port group (MetroCluster 1, MetroCluster 2). You'll need it when using the RcfFileGenerator tool as described later in this configuration procedure.

- The RcfFileGenerator for MetroCluster IP also provides a per-port cabling overview for each switch.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS500f AFF C250, ASA C250 AFF A250, ASA A250	Cisco 9336C-FX2 12-port platform port assignments (group 1)
AFF A20	Cisco 9336C-FX2 12-port platform port assignments (group 2)
AFF A30, AFF C30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 3a) or 100G (group 3b) Ethernet card. <ul style="list-style-type: none">• Cisco 9336C-FX2 12-port platform port assignments (group 3a - 25G)• Cisco 9336C-FX2 12-port platform port assignments (group 3b - 100G)
FAS8300, AFF C400, ASA C400, FAS8700 AFF A400, ASA A400	Cisco 9336C-FX2 12-port platform port assignments (group 4)
AFF A50	Cisco 9336C-FX2 12-port platform port assignments (group 5)
AFF C800, ASA C800, AFF A800, ASA A800 FAS9500, AFF A900, ASA A900	Cisco 9336C-FX2 12-port platform port assignments (group 6)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	Cisco 9336C-FX2 12-port platform port assignments (group 7)

Cisco 9336C-FX2 12-port platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF A150 ASA A150		FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1-4	Unused	disabled		disabled	
5-6	Ports disallowed to use	blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
9/2-4		disabled		disabled	
10/1		e0a	e0b	e0c	e0d
10/2-4		disabled		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b	e0c	e0d
11/2-4		disabled		disabled	
12/1		e0a	e0b	e0c	e0d
12/2-4		disabled		disabled	
13-18	Ports disallowed to use	blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
22/1-4					
23-36	Ports disallowed to use	blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 2)

Review the platform port assignments to cable an AFF A20 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e2a	e4a
1/2-4		disabled	
2/1		e2a	e4a
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e2a	e4a
3/2-4		disabled	
4/1		e2a	e4a
4/2-4		disabled	
5-6	Ports disallowed to use	blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e2b	e4b
9/2-4		disabled	
10/1		e2b	e4b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e2b	e4b
11/2-4		disabled	
12/1		e2b	e4b
12/2-4		disabled	
13-18	Ports disallowed to use	blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster	
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster	
22/1-4			
23-36	Ports disallowed to use	blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 3a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a 12-port Cisco 9336C-FX2 switch using a four-port 25G Ethernet card.



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled		disabled	
2/1		e4a	e4b	e4a	e4b	e4a	e4b
2/2-4		disabled		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled		disabled	
4/1		e4a	e4b	e4a	e4b	e4a	e4b
4/2-4		disabled		disabled		disabled	
5-6	Ports disallowed to use	blocked		blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13-18	Ports disallowed to use	blocked		blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23-36	Ports disallowed to use	blocked		blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 3b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a 12-port Cisco 9336C-FX2 switch using a two-port 100G Ethernet card.



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5-6	Ports disallowed to use	blocked		blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8		ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b	e2a	e2b
13-18	Ports disallowed to use	blocked		blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
23-36	Ports disallowed to use	blocked		blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 4)

Review the platform port assignments to cable an FAS8300, AFF C400, ASA C400, FAS8700, AFF A400, or ASA A400 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0c	e0d	e3a	e3b
2					
3	MetroCluster 2, Local Cluster interface	e0c	e0d	e3a	e3b
4					
5-6	Ports disallowed to use	blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e1a	e1b	e1a	e1b
10					
11	MetroCluster 2, MetroCluster interface	e1a	e1b	e1a	e1b
12					
13-18	Ports disallowed to use	blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
22/1-4					
23-36	Ports disallowed to use	blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 5)

Review the platform port assignments to cable an AFF A50 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2		e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4		e4a	e4b
5-6	Ports disallowed to use	blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10		e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12		e2a	e2b
13-18	Ports disallowed to use	blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster	
20			
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster	
22/1-4			
23-36	Ports disallowed to use	blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Cisco 9336C-FX2 12-port platform port assignments (group 6)

Review the platform port assignments to cable an AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	AFF C800 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a (note 2)
2					
3	MetroCluster 2, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a (note 2)
4					
5-6	Ports disallowed to use	blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e0b	e1b	e5b	e7b
10					
11	MetroCluster 2, MetroCluster interface	e0b	e1b	e5b	e7b
12					
13-18	Ports disallowed to use	blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
20					
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster	
22/1-4					
23-36	Ports disallowed to use	blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Note 2: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

Cisco 9336C-FX2 12-port platform port assignments (group 7)

Review the platform port assignments to cable an AFF A70, FAS70, AFF C80, FAS90, AFF A90, or AFF A1K system to a 12-port Cisco 9336C-FX2 switch:

Switch Port	Port use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5-6	Ports disallowed to use	blocked		blocked		blocked		blocked	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
8									
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
12									
13-18	Ports disallowed to use	blocked		blocked		blocked		blocked	
19	ISL, MetroCluster native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
20									
21/1-4	ISL, MetroCluster breakout mode 10G / 25G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
22/1-4									
23-36	Ports disallowed to use	blocked		blocked		blocked		blocked	

Note 1: You can only configure ports 19 and 20 **or** ports 21 and 22. If you use ports 19 and 20 first, then ports 21 and 22 are blocked. If you use ports 21 and 22 first, then ports 19 and 20 are blocked.

Platform port assignments for a 36-port Cisco 9336C-FX2 switch connecting NS224 storage in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- The tables in this section are for 36-port Cisco 9336C-FX2 switches when at least one MetroCluster configuration or DR group is connecting NS224 shelves to the MetroCluster switch.

If you have a 36-port Cisco 9336C-FX2 switch and you don't plan to connect NS224 storage to the switch, use the tables in [Platform port assignments for Cisco 3232C or 36-port Cisco 9336C-FX2 switches](#).

If you have a 12-port Cisco 9336C-FX2 switch, use the tables in [Platform port assignments for 12-port Cisco 9336C-FX2 switches](#).



The 12-port Cisco 9336C-FX2 switch doesn't support connecting NS224 shelves to the MetroCluster switch.

- When you cable a Cisco 9336C-FX2 switch connecting NS224 storage, you can only have a maximum of two MetroCluster configurations or DR groups. At least one MetroCluster configuration or DR group must be connecting NS224 shelves to the MetroCluster switch. If one of your MetroCluster configurations or DR groups is a system that doesn't support NS224 shelves, it can only be connected as the second MetroCluster configuration or DR group.

If your second MetroCluster or DR group doesn't connect NS224 shelves to the MetroCluster switch, follow the [cabling tables for controllers not connecting switch-attached NS224 shelves](#).

- The RcfFileGenerator only shows eligible platforms when the first platform is selected.
- Connecting one eight-node or two four-node MetroCluster configurations requires ONTAP 9.14.1 or later.

Choose the correct cabling table for your configuration

Review the correct port assignments table for your configuration. There are two sets of cabling tables in this section:

- [Cabling tables for controllers connecting switch-attached NS224 shelves](#)
- [Cabling tables for controllers not connecting switch-attached NS224 shelves](#)

Controllers connecting switch-attached NS224 shelves

Determine which port assignments table you should follow for controllers connecting switch-attached NS224 shelves.

Platform	Use this cabling table...
AFF C30, AFF A30 AFF C60	The table you follow depends on whether you are using a 25G (group 1a) or 100G (group 1b) Ethernet card. <ul style="list-style-type: none"> • Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1a - 25G) • Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1b - 100G)
AFF A320 AFF C400, ASA C400 AFF A400, ASA A400	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 2)
AFF A50	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 3)
AFF A700 AFF C800, ASA C800, AFF A800 AFF A900, ASA A900	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 4)
AFF A70 AFF C80 AFF A90 AFF A1K	Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 5)

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1a)

Review the platform port assignments to cable an AFF A30, AFF C30, or AFF C60 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a four-port 25G Ethernet card.



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers connecting switch-attached shelves					
Switch Port	Port Use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled	
2/1		e4a	e4b	e4a	e4b
2/2-4		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled	
4/1		e4a	e4b	e4a	e4b
4/2-4		disabled		disabled	
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster	
14					
15					
16					
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b	e3a	e3b
18					
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b	e3a	e3b
20					
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)				
24					
25	Storage shelf 4 (6)				
26					
27	Storage shelf 5 (5)				
28					
29	Storage shelf 6 (4)				
30					
31	Storage shelf 7 (3)				
32					
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 1b)

Review the platform port assignments to cable an AFF A30, AFF C30, or AFF C60 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a two-port 100G Ethernet card.



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers connecting switch-attached shelves					
Switch Port	Port Use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
8					
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b
10		e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b
12		e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster	
14					
15					
16					
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b	e3a	e3b
18		e3a	e3b	e3a	e3b
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b	e3a	e3b
20		e3a	e3b	e3a	e3b
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)				
24					
25	Storage shelf 4 (6)				
26					
27	Storage shelf 5 (5)				
28					
29	Storage shelf 6 (4)				
30					
31	Storage shelf 7 (3)				
32					
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 2)

Review the platform port assignments to cable an AFF A320, AFF C400, ASA C400, AFF A400, or ASA A400 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves							
Switch Port	Port Use	AFF A320		AFF C400 ASA C400		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
2							
3	MetroCluster 2, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
4							
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
10							
11	MetroCluster 2, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17	MetroCluster 1, Ethernet Storage Interface	e0c	e0f	e4a	e4b / e5b	e0c	e0d / e5b
18							
19	MetroCluster 2, Ethernet Storage Interface	e0c	e0f	e4a	e4b / e5b	e0c	e0d / e5b
20							
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
24		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
26		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
27	Storage shelf 5 (5)	NSM-1, e0a	NSM-1, e0b				
28		NSM-2, e0a	NSM-2, e0b				
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
30		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
32		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 3)

Review the platform port assignments to cable an AFF A50 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves			
Switch Port	Port Use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2		e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4		e4a	e4b
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10		e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12		e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b
18			
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b
20			
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b
24		NSM-2, e0a	NSM-2, e0b
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b
26		NSM-2, e0a	NSM-2, e0b
27	Storage shelf 5 (5)		
28			
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b
30		NSM-2, e0a	NSM-2, e0b
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b
32		NSM-2, e0a	NSM-2, e0b
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 4)

Review the platform port assignments to cable an AFF A700, AFF C800, ASA C800, AFF A800, AFF A900, or ASA A900 system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves							
Switch Port	Port Use	AFF A700		AFF C800 ASA C800 AFF A800		AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
2							
3	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
4							
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
6		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
10							
11	MetroCluster 2, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17	MetroCluster 1, Ethernet Storage Interface	e3a	e3b / e7b	e5a	e5b / e3b	e3a (option 1) e2a (option 2)	e3b (option 1) e10b (option 2)
18							
19	MetroCluster 2, Ethernet Storage Interface	e3a	e3b / e7b	e5a	e5b / e3b	e3a (option 1) e2a (option 2)	e3b (option 1) e10b (option 2)
20							
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
22		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
24		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
26		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
27	Storage shelf 5 (5)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
28		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
30		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
32		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
34		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b	NSM-1, e0a	NSM-1, e0b
36		NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b	NSM-2, e0a	NSM-2, e0b

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

Cisco 9336C-FX2 switch connecting NS224 storage platform port assignments (group 5)

Review the platform port assignments to cable an AFF A70, AFF C80, AFF A90, or AFF A1K system that is connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers connecting switch-attached shelves									
Switch Port	Port Use	AFF A70		AFF C80		AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5	Storage shelf 1 (9)	NSM-1, e0a	NSM-1, e0b						
6		NSM-2, e0a	NSM-2, e0b						
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
8									
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
12									
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14									
15									
16									
17	MetroCluster 1, Ethernet Storage Interface	e8a (option 1)	e8b (option 1)						
18		e11a (option 2)	e11b (option 2)						
19	MetroCluster 2, Ethernet Storage Interface	e8a (option 1)	e8b (option 1)						
20		e11a (option 2)	e11b (option 2)						
21	Storage shelf 2 (8)	NSM-1, e0a	NSM-1, e0b						
22		NSM-2, e0a	NSM-2, e0b						
23	Storage shelf 3 (7)	NSM-1, e0a	NSM-1, e0b						
24		NSM-2, e0a	NSM-2, e0b						
25	Storage shelf 4 (6)	NSM-1, e0a	NSM-1, e0b						
26		NSM-2, e0a	NSM-2, e0b						
27	Storage shelf 5 (5)	NSM-1, e0a	NSM-1, e0b						
28		NSM-2, e0a	NSM-2, e0b						
29	Storage shelf 6 (4)	NSM-1, e0a	NSM-1, e0b						
30		NSM-2, e0a	NSM-2, e0b						
31	Storage shelf 7 (3)	NSM-1, e0a	NSM-1, e0b						
32		NSM-2, e0a	NSM-2, e0b						
33	Storage shelf 8 (2)	NSM-1, e0a	NSM-1, e0b						
34		NSM-2, e0a	NSM-2, e0b						
35	Storage shelf 9 (1)	NSM-1, e0a	NSM-1, e0b						
36		NSM-2, e0a	NSM-2, e0b						

Controllers not connecting switch-attached NS224 shelves

Determine which port assignments table you should follow for controllers that are not connecting switch-attached NS224 shelves.

Platform	Use this cabling table...
AFF A150, ASAA150 FAS2750, AFF A220	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 6)
AFF A20	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 7)
FAS500f AFF C250, ASA C250 AFF A250, ASA A250	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 8)

Platform	Use this cabling table...
AFF C30, AFF A30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 9a) or 100G (group 9b) Ethernet card. <ul style="list-style-type: none"> • Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9a) • Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9b)
FAS8200, AFF A300	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 10)
AFF A320 FAS8300, AFF C400, ASA C400, FAS8700 AFF A400, ASA A400	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 11)
AFF A50	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 12)
FAS9000, AFF A700 AFF C800, ASA C800, AFF A800, ASA A800 FAS9500, AFF A900, ASA A900	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 13)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 14)

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 6)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS2750, or AFF A220 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	AFF A150 ASA A150 FAS2750 AFF A220	
		IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b
9/2-4		disabled	
10/1		e0a	e0b
10/2-4		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b
11/2-4		disabled	
12/1		e0a	e0b
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 7)

Review the platform port assignments to cable an AFF A20 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e2a	e4a
1/2-4		disabled	
2/1		e2a	e4a
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e2a	e4a
3/2-4		disabled	
4/1		e2a	e4a
4/2-4		disabled	
5-6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e2b	e4b
9/2-4		disabled	
10/1		e2b	e4b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e2b	e4b
11/2-4		disabled	
12/1		e2b	e4b
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 8)

Review the platform port assignments to cable a FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0c	e0d
9/2-4		disabled	
10/1		e0c	e0d
10/2-4		disabled	
11/1	MetroCluster 2, Shared Cluster and MetroCluster interface	e0c	e0d
11/2-4		disabled	
12/1		e0c	e0d
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a four-port 25G Ethernet card:



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers not connecting switch-attached shelves							
Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
1/2-4		disabled		disabled		disabled	
2/1		e4a	e4b	e4a	e4b	e4a	e4b
2/2-4		disabled		disabled		disabled	
3/1	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3/2-4		disabled		disabled		disabled	
4/1		e4a	e4b	e4a	e4b	e4a	e4b
4/2-4		disabled		disabled		disabled	
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17-36	Unused	disabled		disabled		disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 9b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch using a two-port 100G Ethernet card:



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Controllers not connecting switch-attached shelves							
Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2		e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4		e4a	e4b	e4a	e4b	e4a	e4b
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17-36	Unused	disabled		disabled		disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 10)

Review the platform port assignments to cable a FAS8200 or AFF A300 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port Use	FAS8200 AFF A300	
		IP_Switch_x_1	IP_Switch_x_2
1/1	MetroCluster 1, Local Cluster interface	e0a	e0b
1/2-4		disabled	
2/1		e0a	e0b
2/2-4		disabled	
3/1	MetroCluster 2, Local Cluster interface	e0a	e0b
3/2-4		disabled	
4/1		e0a	e0b
4/2-4		disabled	
5-6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9/1	MetroCluster 1, MetroCluster interface	e1a	e1b
9/2-4		disabled	
10/1		e1a	e1b
10/2-4		disabled	
11/1	MetroCluster 2, MetroCluster interface	e1a	e1b
11/2-4		disabled	
12/1		e1a	e1b
12/2-4		disabled	
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 11)

Review the platform port assignments to cable an AFF A320, FAS8300, AFF C400, ASA C400, FAS8700, AFF A400, or ASA A400 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves							
Switch Port	Port Use	AFF A320		FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
2							
3	MetroCluster 2, Local Cluster interface	e0a	e0d	e0c	e0d	e3a	e3b
4							
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
10							
11	MetroCluster 2, MetroCluster interface	e0g	e0h	e1a	e1b	e1a	e1b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17-36	Unused	disabled		disabled		disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 12)

Review the platform port assignments to cable an AFF A50 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves			
Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2			
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4			
5-6	Unused	disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8			
9	MetroCluster 1, MetroCluster interface	e2a	e2b
10			
11	MetroCluster 2, MetroCluster interface	e2a	e2b
12			
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17-36	Unused	disabled	

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 13)

Review the platform port assignments to cable a FAS9000, AFF A700, AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves							
Switch Port	Port Use	FAS9000 AFF A700		AFF C800 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
2							
3	MetroCluster 2, Local Cluster interface	e4a	e4e / e8a	e0a	e1a	e4a	e4b(e) / e8a Note 1
4							
5-6	Unused	disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
8							
9	MetroCluster 1, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
10							
11	MetroCluster 2, MetroCluster interface	e5a	e5b	e0b	e1b	e5b	e7b
12							
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17-36	Unused	disabled		disabled		disabled	

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

Cisco 9336C-FX2 switch not connecting NS224 storage platform port assignments (group 14)

Review the platform port assignments to cable an AFF A70, FAS70, AFF C80, FAS90, AFF A90, or AFF A1K system that isn't connecting switch-attached NSS24 shelves to a Cisco 9336C-FX2 switch:

Controllers not connecting switch-attached shelves									
Switch Port	Port Use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5-6	Unused	disabled		disabled		disabled		disabled	
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster							
8									
9	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2b	e3b	e2b	e3b
12									
13	ISL MetroCluster, native speed 40G / 100G breakout mode 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14									
15									
16									
17-36	Unused	disabled		disabled		disabled		disabled	

Platform port assignments for Broadcom supported BES-53248 IP switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- You cannot use the switches with remote ISL ports of different speeds (for example, a 25 Gbps port connected to a 10 Gbps ISL port).
- If you configure the switch for MetroCluster FC to IP Transition, the following ports are used depending on the target platform that you choose:

Target platform	Port
FAS500f, AFF C250, ASA C250, AFF A250, ASA A250, FAS8300, AFF C400, ASA C400, AFF A400, ASA A400, or FAS8700 platforms	ports 1 - 6, 10Gbps
FAS8200 or AFF A300 platforms	ports 3 - 4 and 9 - 12, 10Gbps

- AFF A320 systems configured with Broadcom BES-53248 switches might not support all features.

Any configuration or feature that requires that the local cluster connections are connected to a switch is not supported. For example, the following configurations and procedures are not supported:

- Eight-node MetroCluster configurations
- Transitioning from MetroCluster FC to MetroCluster IP configurations
- Refreshing a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS2750 AFF A220	Broadcom BES-53248 platform port assignments (group 1)
FAS500f AFF C250, ASA C250 AFF A250, ASA A250	Broadcom BES-53248 platform port assignments (group 2)
AFF A20	Broadcom BES-53248 platform port assignments (group 3)
AFF C30, AFF A30 FAS50 AFF C60	Broadcom BES-53248 platform port assignments (group 4)
FAS8200, AFF A300	Broadcom BES-53248 platform port assignments (group 5)
AFF A320	Broadcom BES-53248 platform port assignments (group 6)
FAS8300 AFF C400, ASA C400 AFF A400, ASA A400 FAS8700	Broadcom BES-53248 platform port assignments (group 7)

Broadcom BES-53248 platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS2750, or AFF A220 system to a Broadcom BES-53248 switch:

Physical Port	Port use	AFF A150 ASA A150 FAS2750 AFF A220	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Shared Cluster and MetroCluster interface	e0a	e0b
2			
3	MetroCluster 2, Shared Cluster and MetroCluster interface	e0a	e0b
4			
5-8	Unused	disabled	
9	MetroCluster 3, Shared Cluster and MetroCluster interface	e0a	e0b
10			
11	MetroCluster 4, Shared Cluster and MetroCluster interface	e0a	e0b
12			
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (Note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.
- If both MetroCluster configurations are using the same the platform, NetApp recommends selecting group "MetroCluster 3" for one configuration and group "MetroCluster 4" for the other configuration. If the platforms are different, then you must select "MetroCluster 3" or "MetroCluster 4" for the first configuration, and "MetroCluster 1" or "MetroCluster 2" for the second configuration.

Broadcom BES-53248 platform port assignments (group 2)

Review the platform port assignments to cable a FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a Broadcom BES-53248 switch:

Physical Port	Port use	FAS500f AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2
1 - 4	Unused	disabled	
5	MetroCluster 1, Shared Cluster and MetroCluster interface	e0c	e0d
6			
7	MetroCluster 2, Shared Cluster and MetroCluster interface	e0c	e0d
8			
9	MetroCluster 3, Shared Cluster and MetroCluster interface	e0c	e0d
10			
11	MetroCluster 4, Shared Cluster and MetroCluster interface	e0c	e0d
12			
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (Note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.
- If both MetroCluster configurations are using the same the platform, NetApp recommends selecting group "MetroCluster 3" for one configuration and group "MetroCluster 4" for the other configuration. If the platforms are different, then you must select "MetroCluster 3" or "MetroCluster 4" for the first configuration, and "MetroCluster 1" or "MetroCluster 2" for the second configuration.

Broadcom BES-53248 platform port assignments (group 3)

Review the platform port assignments to cable an AFF A20 system to a Broadcom BES-53248 switch:

Physical Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e2a	e4a
2			
3	MetroCluster 2, Local Cluster interface	e2a	e4a
4			
5	MetroCluster 1, MetroCluster interface	e2b	e4b
6			
7	MetroCluster 2, MetroCluster interface	e2b	e4b
8			
9 - 12	Unused	disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
17	MetroCluster 3, Local Cluster interface (note 1)	e2a	e4a
18			
19	MetroCluster 3, MetroCluster interface (note 1)	e2b	e4b
20			
21	MetroCluster 4, Local Cluster interface (note 1)	e2a	e4a
22			
23	MetroCluster 4, MetroCluster interface (note 1)	e2b	e4b
24			
..	Ports not licensed (25 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.

Broadcom BES-53248 platform port assignments (group 4)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a Broadcom BES-53248 switch using a four-port 25G Ethernet card.



- This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.
- This configuration requires a QSFP-to-SFP+ adapter in the card on the controller to support a 25Gbps network speed.

Physical Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2							
3	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4							
5	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
6							
7	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
8							
9 - 12	Unused	disabled		disabled		disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14							
15							
16							
17	MetroCluster 3, Local Cluster interface (note 1)	e4a	e4b	e4a	e4b	e4a	e4b
18							
19	MetroCluster 3, MetroCluster interface (note 1)	e2a	e2b	e2a	e2b	e2a	e2b
20							
21	MetroCluster 4, Local Cluster interface (note 1)	e4a	e4b	e4a	e4b	e4a	e4b
22							
23	MetroCluster 4, MetroCluster interface (note 1)	e2a	e2b	e2a	e2b	e2a	e2b
24							
..	Ports not licensed (25 - 54)						
53	ISL, MetroCluster, native speed 40G / 100G (note 1)	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
54							
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
56							

- **Note 1:** Using these ports requires an additional license.

Broadcom BES-53248 platform port assignments (group 5)

Review the platform port assignments to cable a FAS8200 or AFF A300 system to a Broadcom BES-53248 switch:

Physical Port	Port use	FAS8200 AFF A300	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e0b
2			
3	MetroCluster 2, Local Cluster interface	e0a	e0b
4			
5	MetroCluster 1, MetroCluster interface	e1a	e1b
6			
7	MetroCluster 2, MetroCluster interface	e1a	e1b
8			
9 - 12	Unused	disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (note 1)	ISL, MetroCluster	
54			
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
56			

- **Note 1:** Using these ports requires an additional license.

Broadcom BES-53248 platform port assignments (group 6)

Review the platform port assignments to cable an AFF A320 system to a Broadcom BES-53248 switch:

Physical Port	Port use	AFF A320	
		IP_Switch_x_1	IP_Switch_x_2
1 - 12	Ports not used (Note 2)	disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster	
14			
15			
16			
..	Ports not licensed (17 - 54)		
53	ISL, MetroCluster, native speed 40G / 100G (see Note 1)	ISL, MetroCluster	
54			
55	MetroCluster 1, MetroCluster interface (Note 2)	e0g	e0h
56			

- **Note 1:** Using these ports requires an additional license.
- **Note 2:** Only a single four-node MetroCluster using AFF A320 systems can be connected to the switch.

Features that require a switched cluster are not supported in this configuration. This includes the MetroCluster FC to IP transition and tech refresh procedures.

Broadcom BES-53248 platform port assignments (group 7)

Review the platform port assignments to cable a FAS8300, AFF C400, ASA C400, AFF A400, ASA A400, or FAS8700 system to a Broadcom BES-53248 switch:

Physical Port	Port use	FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1 - 12	Ports not used (see Note 2)	disabled		disabled	
13	ISL, MetroCluster native speed 10G / 25G	ISL, MetroCluster		ISL, MetroCluster	
14					
15					
16					
..	Ports not licensed (17 - 48)				
49	MetroCluster 5, Local Cluster interface (Note 1)	e0c	e0d	e3a	e3b
50					
51	MetroCluster 5, MetroCluster interface (Note 1)	e1a	e1b	e1a	e1b
52					
53	ISL, MetroCluster, native speed 40G / 100G (Note 1)	ISL, MetroCluster		ISL, MetroCluster	
54					
55	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster	
56					

- **Note 1:** Using these ports requires an additional license.

- **Note 2:** Only a single four-node MetroCluster using AFF A320 systems can be connected to the switch.

Features that require a switched cluster are not supported in this configuration. This includes the MetroCluster FC to IP transition and tech refresh procedures.

Platform port assignments for NVIDIA supported SN2100 IP switches in a MetroCluster IP configuration

The port usage in a MetroCluster IP configuration depends on the switch model and platform type.

Review the following considerations before using the configuration tables:

- Connecting an eight-node or two four-node MetroCluster configurations requires ONTAP 9.14.1 or later and RCF file version 2.00 or later.



The RCF file version is different to the version of the RCFfilegenerator tool used to generate the file. For example, you can generate an RCF file version 2.00 using RCFfilegenerator v1.6c.

- If you cable multiple MetroCluster configurations then follow the respective table.
For example:
 - If you cable two four-node MetroCluster configurations of type AFF A700, then connect the first MetroCluster shown as "MetroCluster 1", and the second MetroCluster shown as "MetroCluster 2" in the AFF A700 table.



Ports 13 and 14 can be used in native speed mode supporting 40 Gbps and 100 Gbps, or in breakout mode to support 4 × 25 Gbps or 4 × 10 Gbps. If they use native speed mode they are represented as ports 13 and 14. If they use breakout mode, either 4 × 25 Gbps or 4 × 10 Gbps, then they are represented as ports 13s0-3 and 14s0-3.

The following sections describe the physical cabling outline. You can also refer to the [RcfFileGenerator](#) for detailed cabling information.

Choose the correct cabling table for your configuration

Use the following table to determine which cabling table you should follow.

If your system is...	Use this cabling table...
AFF A150, ASA A150 FAS500f AFF C250, ASA C250 AFF A250, ASA A250	NVIDIA SN2100 platform port assignments (group 1)
AFF A20	NVIDIA SN2100 platform port assignments (group 2)

If your system is...	Use this cabling table...
AFF C30, AFF A30 FAS50 AFF C60	The table you follow depends on whether you are using a 25G (group 3a) or 100G (group 3b) Ethernet card. <ul style="list-style-type: none"> • NVIDIA SN2100 platform port assignments (group 3a -25G) • NVIDIA SN2100 platform port assignments (group 3b -100G)
FAS8300 AFF C400, ASA C400 AFF A400, ASA A400 FAS8700 FAS9000, AFF A700	NVIDIA SN2100 platform port assignments (group 4)
AFF A50	NVIDIA SN2100 platform port assignments (group 5)
AFF C800, ASA C800 AFF A800, ASA A800 FAS9500 AFF A900, ASA A900	NVIDIA SN2100 platform port assignments (group 6)
FAS70, AFF A70 AFF C80 FAS90, AFF A90 AFF A1K	NVIDIA SN2100 platform port assignments (group 7)

NVIDIA SN2100 platform port assignments (group 1)

Review the platform port assignments to cable an AFF A150, ASA A150, FAS500f, AFF C250, ASA C250, AFF A250, or ASA A250 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF A150 ASA A150		FAS500F AFF C250 ASA C250 AFF A250 ASA A250	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1 - 6	Unused	disabled		disabled	
7s0	MetroCluster 1, Shared Cluster and MetroCluster interface	e0c	e0d	e0c	e0d
7s1-3		disabled		disabled	
8s0		e0c	e0d	e0c	e0d
8s1-3		disabled		disabled	
9s0	MetroCluster 2, Shared Cluster and MetroCluster interface	e0c	e0d	e0c	e0d
9s1-3		disabled		disabled	
10s0		e0c	e0d	e0c	e0d
10s1-3		disabled		disabled	
11s0	MetroCluster 3, Shared Cluster and MetroCluster interface	e0c	e0d	e0c	e0d
11s1-3		disabled		disabled	
12s0		e0c	e0d	e0c	e0d
12s1-3		disabled		disabled	
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3		ISL, MetroCluster		ISL, MetroCluster	
15	ISL, Local Cluster	ISL, Local Cluster		ISL, Local Cluster	
16	100G	ISL, Local Cluster		ISL, Local Cluster	

NVIDIA SN2100 platform port assignments (group 2)

Review the platform port assignments to cable an AFF A20 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF A20	
		IP_Switch_x_1	IP_Switch_x_2
1s0	MetroCluster 1, Local Cluster interface	e2a	e4a
s1s1-3		disabled	
2s0		e2a	e4a
2s1-3		disabled	
3s0	MetroCluster 2, Local Cluster interface	e2a	e4a
3s1-3		disabled	
4s0		e2a	e4a
4s1-3		disabled	
5s0	MetroCluster 3, Local Cluster interface	e2a	e4a
5s1-3		disabled	
6s0		e2a	e4a
6s1-3		disabled	
7	MetroCluster 1, MetroCluster interface	e2b	e4b
8			
9	MetroCluster 2, MetroCluster interface	e2b	e4b
10			
11	MetroCluster 3, MetroCluster interface	e2b	e4b
12			
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster	
14 / 14s0-3			
15	ISL, Local Cluster 100G	ISL, Local Cluster	
16			

NVIDIA SN2100 platform port assignments (group 3a)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a NVIDIA SN2100 switch using a four-port 25G Ethernet card:



This configuration requires a four-port 25G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (25G Cluster/HA) AFF A30 (25G Cluster/HA)		FAS50 (25G Cluster/HA)		AFF C60 (25G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1s0	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
s1s1-3		disabled		disabled		disabled	
2s0		e4a	e4b	e4a	e4b	e4a	e4b
2s1-3		disabled		disabled		disabled	
3s0	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3s1-3		disabled		disabled		disabled	
4s0		e4a	e4b	e4a	e4b	e4a	e4b
4s1-3		disabled		disabled		disabled	
5s0	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
5s1-3		disabled		disabled		disabled	
6s0		e4a	e4b	e4a	e4b	e4a	e4b
6s1-3		disabled		disabled		disabled	
7	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
8							
9	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
16	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	

NVIDIA SN2100 platform port assignments (group 3b)

Review the platform port assignments to cable an AFF A30, AFF C30, AFF C60, or FAS50 system to a NVIDIA SN2100 switch using a two-port 100G Ethernet card:



This configuration requires a two-port 100G Ethernet card in slot 4 to connect the local cluster and HA interfaces.

Switch Port	Port use	AFF C30 (100G Cluster/HA) AFF A30 (100G Cluster/HA)		FAS50 (100G Cluster/HA)		AFF C60 (100G Cluster/HA)	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
2	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
3	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
4	MetroCluster 1, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
5	MetroCluster 2, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
6	MetroCluster 3, Local Cluster interface	e4a	e4b	e4a	e4b	e4a	e4b
7	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
8	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
9	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
10	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
11	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
12	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e2b	e2a	e2b
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
16	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	

NVIDIA SN2100 platform port assignments (group 4)

Review the platform port assignments to cable a FAS8300, AFF C400, ASA C400, AFF A400, ASA A400, FAS8700, FAS9000, or AFF A700 system to a NVIDIA SN2100 switch:

Switch Port	Port use	FAS8300 AFF C400 ASA C400 FAS8700		AFF A400 ASA A400		FAS9000 AFF A700	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0c	e0d	e3a	e3b	e4a	e4e / e8a Note 1
2							
3	MetroCluster 2, Local Cluster interface	e0c	e0d	e3a	e3b	e4a	e4e / e8a Note 1
4							
5	MetroCluster 3, Local Cluster interface	e0c	e0d	e3a	e3b	e4a	e4e / e8a Note 1
6							
7	MetroCluster 1, MetroCluster interface	e1a	e1b	e1a	e1b	e5a	e5b
8							
9	MetroCluster 2, MetroCluster interface	e1a	e1b	e1a	e1b	e5a	e5b
10							
11	MetroCluster 3, MetroCluster interface	e1a	e1b	e1a	e1b	e5a	e5b
12							
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3							
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster		ISL, Local Cluster	
16							

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

NVIDIA SN2100 platform port assignments (group 5)

Review the platform port assignments to cable an AFF A50 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF A50	
		IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e4a	e4b
2			
3	MetroCluster 2, Local Cluster interface	e4a	e4b
4			
5	MetroCluster 3, Local Cluster interface	e4a	e4b
6			
7	MetroCluster 1, MetroCluster interface	e2a	e2b
8			
9	MetroCluster 2, MetroCluster interface	e2a	e2b
10			
11	MetroCluster 3, MetroCluster interface	e2a	e2b
12			
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster	
14 / 14s0-3			
15	ISL, Local Cluster 100G	ISL, Local Cluster	
16			

NVIDIA SN2100 platform port assignments (group 6)

Review the platform port assignments to cable an AFF C800, ASA C800, AFF A800, ASA A800, FAS9500, AFF A900, or ASA A900 system to a NVIDIA SN2100 switch:

Switch Port	Port use	AFF C80 ASA C800 AFF A800 ASA A800		FAS9500 AFF A900 ASA A900	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a Note 1
2					
3	MetroCluster 2, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a Note 1
4					
5	MetroCluster 3, Local Cluster interface	e0a	e1a	e4a	e4b(e) / e8a Note 1
6					
7	MetroCluster 1, MetroCluster interface	e0b	e1b	e5b	e7b
8					
9	MetroCluster 2, MetroCluster interface	e0b	e1b	e5b	e7b
10					
11	MetroCluster 3, MetroCluster interface	e0b	e1b	e5b	e7b
12					
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3					
15	ISL, Local Cluster 100G	ISL, Local Cluster		ISL, Local Cluster	
16					

Note 1: Use either ports e4a and e4e or e4a and e8a if you are using an X91440A adapter (40Gbps). Use either ports e4a and e4b or e4a and e8a if you are using an X91153A adapter (100Gbps).

NVIDIA SN2100 platform port assignments (group 7)

Review the platform port assignments to cable a FAS70, AFF A70, AFF C80, FAS90, AFF A90, or AFF A1K system to a NVIDIA SN2100 switch:

Switch Port	Port use	FAS70 AFF A70		AFF C80		FAS90 AFF A90		AFF A1K	
		IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2	IP_Switch_x_1	IP_Switch_x_2
1	MetroCluster 1, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
2									
3	MetroCluster 2, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
4									
5	MetroCluster 3, Local Cluster interface	e1a	e7a	e1a	e7a	e1a	e7a	e1a	e7a
6									
7	MetroCluster 1, MetroCluster interface	e2a	e2b	e2a	e3b	e2b	e3b	e2b	e3b
8									
9	MetroCluster 2, MetroCluster interface	e2a	e2b	e2a	e3b	e2b	e3b	e2b	e3b
10									
11	MetroCluster 3, MetroCluster interface	e2a	e2b	e2a	e3b	e2b	e3b	e2b	e3b
12									
13 / 13s0-3	MetroCluster ISL 40/100G or 4x25G or 4x10G	ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster		ISL, MetroCluster	
14 / 14s0-3									
15	ISL, Local Cluster 100G	ISL, Local Cluster							
16									

Cable the ONTAP controller module ports in a MetroCluster IP configuration

You must cable the controller module ports used for cluster peering, management and data connectivity.

This task must be performed on each controller module in the MetroCluster configuration.

At least two ports on each controller module should be used for cluster peering.

The recommended minimum bandwidth for the ports and network connectivity is 1 GbE.

1. Identify and cable at least two ports for cluster peering and verify they have network connectivity with the partner cluster.

Cluster peering can be done on dedicated ports or on data ports. Using dedicated ports provides higher throughput for the cluster peering traffic.

[Cluster and SVM peering express configuration](#)

2. Cable the controller's management and data ports to the management and data networks at the local site.

Use the installation instructions for your platform at the [ONTAP Hardware Systems Documentation](#).



MetroCluster IP systems do not have dedicated high-availability (HA) ports. Depending on your platform, HA traffic is served using the MetroCluster, local cluster, or shared cluster/MetroCluster interface. When using *ONTAP Hardware Systems Documentation* to install your platform, you should not follow the instructions to cable the cluster and HA ports.

Configure the MetroCluster IP switches

Choose the correct MetroCluster IP switch configuration procedure

You must configure the IP switches to provide backend MetroCluster IP connectivity. The procedure you follow depends on your switch vendor.

- [Configure Broadcom IP switches](#)
- [Configure Cisco IP switches](#)
- [Configure NVIDIA IP switches](#)

Configure Broadcom IP switches for cluster interconnect and backend MetroCluster IP connectivity

You must configure the Broadcom IP switches for use as the cluster interconnect and for backend MetroCluster IP connectivity.



Your configuration requires additional licenses (6 x 100-Gb port license) in the following scenarios:

- You use ports 53 and 54 as a 40-Gbps or 100-Gbps MetroCluster ISL.
- You use a platform that connects the local cluster and MetroCluster interfaces to ports 49 - 52.

Resetting the Broadcom IP switch to factory defaults

Before installing a new switch software version and RCFs, you must erase the Broadcom switch settings and perform basic configuration.

About this task

- You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.
- You must be connected to the switch using the serial console.
- This task resets the configuration of the management network.

Steps

1. Change to the elevated command prompt (#): `enable`

```
(IP_switch_A_1)> enable
(IP_switch_A_1) #
```

2. Erase the startup configuration and remove the banner

a. Erase the startup configuration:

erase startup-config

```
(IP_switch_A_1) #erase startup-config

Are you sure you want to clear the configuration? (y/n) y

(IP_switch_A_1) #
```

This command does not erase the banner.

b. Remove the banner:

no set clibanner

```
(IP_switch_A_1) #configure
(IP_switch_A_1) (Config) # no set clibanner
(IP_switch_A_1) (Config) #
```

3. Reboot the switch:

(IP_switch_A_1) #reload

```
Are you sure you would like to reset the system? (y/n) y
```



If the system asks whether to save the unsaved or changed configuration before reloading the switch, select **No**.

4. Wait for the switch to reload, and then log in to the switch.

The default user is “admin”, and no password is set. A prompt similar to the following is displayed:

```
(Routing) >
```

5. Change to the elevated command prompt:

```
enable
```

```
Routing)> enable
(Routing) #
```

6. Set the service port protocol to none:

```
serviceport protocol none
```

```
(Routing) #serviceport protocol none
Changing protocol mode will reset ip configuration.
Are you sure you want to continue? (y/n) y

(Routing) #
```

7. Assign the IP address to the service port:

```
serviceport ip ip-address netmask gateway
```

The following example shows a service port assigned IP address "10.10.10.10" with subnet "255.255.255.0" and gateway "10.10.10.1":

```
(Routing) #serviceport ip 10.10.10.10 255.255.255.0 10.10.10.1
```

8. Verify that the service port is correctly configured:

```
show serviceport
```

The following example shows that the port is up and the correct addresses have been assigned:

```
(Routing) #show serviceport

Interface Status..... Up
IP Address..... 10.10.10.10
Subnet Mask..... 255.255.255.0
Default Gateway..... 10.10.10.1
IPv6 Administrative Mode..... Enabled
IPv6 Prefix is .....
fe80::dac4:97ff:fe56:87d7/64
IPv6 Default Router..... fe80::222:bdff:fef8:19ff
Configured IPv4 Protocol..... None
Configured IPv6 Protocol..... None
IPv6 AutoConfig Mode..... Disabled
Burned In MAC Address..... D8:C4:97:56:87:D7

(Routing) #
```

9. Configure the SSH server.



- The RCF file disables the Telnet protocol. If you do not configure the SSH server, you can only access the bridge using the serial port connection.
- You must configure the SSH server in order to use log collection and other external tools.

a. Generate RSA keys.

```
(Routing) #configure
(Routing) (Config)#crypto key generate rsa
```

b. Generate DSA keys (optional)

```
(Routing) #configure
(Routing) (Config)#crypto key generate dsa
```

c. If you are using the FIPS compliant version of EFOS, generate the ECDSA keys. The following example creates the keys with a length of 521. Valid values are 256, 384 or 521.

```
(Routing) #configure
(Routing) (Config)#crypto key generate ecdsa 521
```

d. Enable the SSH server.

If necessary, exit the configuration context.

```
(Routing) (Config)#end
(Routing) #ip ssh server enable
```



If keys already exist, then you might be asked to overwrite them.

10. If desired, configure the domain and name server:

```
configure
```

The following example shows the `ip domain` and `ip name server` commands:

```
(Routing) # configure
(Routing) (Config)#ip domain name lab.netapp.com
(Routing) (Config)#ip name server 10.99.99.1 10.99.99.2
(Routing) (Config)#exit
(Routing) (Config)#
```

11. If desired, configure the time zone and time synchronization (SNTP).

The following example shows the `sntp` commands, specifying the IP address of the SNTP server and the relative time zone.

```
(Routing) #
(Routing) (Config)#sntp client mode unicast
(Routing) (Config)#sntp server 10.99.99.5
(Routing) (Config)#clock timezone -7
(Routing) (Config)#exit
(Routing) (Config)#
```

For EFOS version 3.10.0.3 and later, use the `ntp` command, as shown in the following example:

```

> (Config)# ntp ?

authenticate          Enables NTP authentication.
authentication-key     Configure NTP authentication key.
broadcast             Enables NTP broadcast mode.
broadcastdelay        Configure NTP broadcast delay in microseconds.
server               Configure NTP server.
source-interface      Configure the NTP source-interface.
trusted-key           Configure NTP authentication key number for
trusted time source.
vrf                   Configure the NTP VRF.

>(Config)# ntp server ?

ip-address|ipv6-address|hostname  Enter a valid IPv4/IPv6 address or
hostname.

>(Config)# ntp server 10.99.99.5

```

12. Configure the switch name:

```
hostname IP_switch_A_1
```

The switch prompt will display the new name:

```

(Routing) # hostname IP_switch_A_1

(IP_switch_A_1) #

```

13. Save the configuration:

```
write memory
```

You receive prompts and output similar to the following example:

```
(IP_switch_A_1) #write memory
```

```
This operation may take a few minutes.  
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Config file 'startup-config' created successfully .
```

```
Configuration Saved!
```

```
(IP_switch_A_1) #
```

14. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Downloading and installing the Broadcom switch EFOS software

You must download the switch operating system file and RCF file to each switch in the MetroCluster IP configuration.

About this task

This task must be repeated on each switch in the MetroCluster IP configuration.

Note the following:

- When upgrading from EFOS 3.4.x.x to EFOS 3.7.x.x or later, the switch must be running EFOS 3.4.4.6 (or later 3.4.x.x release). If you are running a release prior to that, then upgrade the switch to EFOS 3.4.4.6 (or later 3.4.x.x release) first, then upgrade the switch to EFOS 3.7.x.x or later.
- The configuration for EFOS 3.4.x.x and 3.7.x.x or later are different. Changing the EFOS version from 3.4.x.x to 3.7.x.x or later, or vice versa, requires the switch to be reset to factory defaults and the RCF files for the corresponding EFOS version to be (re)applied. This procedure requires access through the serial console port.
- Beginning with EFOS version 3.7.x.x or later, a non-FIPS compliant and a FIPS compliant version is available. Different steps apply when moving to from a non-FIPS compliant to a FIPS compliant version or vice versa. Changing EFOS from a non-FIPS compliant to a FIPS compliant version or vice versa will reset the switch to factory defaults. This procedure requires access through the serial console port.

Steps

1. Download the switch firmware from the [Broadcom support site](#).
2. Check if your version of EFOS is FIPS compliant or non-FIPS compliant by using the `show fips status` command. In the following examples, `IP_switch_A_1` is using FIPS compliant EFOS and `IP_switch_A_2` is using non-FIPS compliant EFOS.

Example 1

```
IP_switch_A_1 #show fips status

System running in FIPS mode

IP_switch_A_1 #
```

Example 2

```
IP_switch_A_2 #show fips status
                ^
% Invalid input detected at `^` marker.

IP_switch_A_2 #
```

3. Use the following table to determine which method you must follow:

Procedure	Current EFOS version	New EFOS version	High level steps
-----------	----------------------	------------------	------------------

Steps to upgrade EFOS between two (non) FIPS compliant versions	3.4.x.x	3.4.x.x	Install the new EFOS image using method 1) The configuration and license information is retained
	3.4.4.6 (or later 3.4.x.x)	3.7.x.x or later non-FIPS compliant	Upgrade EFOS using method 1. Reset the switch to factory defaults and apply the RCF file for EFOS 3.7.x.x or later
	3.7.x.x or later non-FIPS compliant	3.4.4.6 (or later 3.4.x.x)	Downgrade EFOS using method 1. Reset the switch to factory defaults and apply the RCF file for EFOS 3.4.x.x
		3.7.x.x or later non-FIPS compliant	Install the new EFOS image using method 1. The configuration and license information is retained
	3.7.x.x or later FIPS compliant	3.7.x.x or later FIPS compliant	Install the new EFOS image using method 1. The configuration and license information is retained
Steps to upgrade to/from a FIPS compliant EFOS version	Non-FIPS compliant	FIPS compliant	Installation of the EFOS image using method 2. The switch configuration and license information will be lost.
	FIPS compliant	Non-FIPS compliant	

- Method 1: [Steps to upgrade EFOS with downloading the software image to the backup boot partition](#)
- Method 2: [Steps to upgrade EFOS using the ONIE OS installation](#)

Steps to upgrade EFOS with downloading the software image to the backup boot partition

You can perform the following steps only if both EFOS versions are non-FIPS compliant or both EFOS versions are FIPS compliant.



Do not use these steps if one version is FIPS compliant and the other version is non-FIPS compliant.

Steps

1. Copy the switch software to the switch: `copy sftp://user@50.50.50.50/switchsoftware/efos-3.4.4.6.stk backup`

In this example, the efos-3.4.4.6.stk operating system file is copied from the SFTP server at 50.50.50.50 to the backup partition. You need to use the IP address of your TFTP/SFTP server and the file name of the RCF file that you need to install.

```
(IP_switch_A_1) #copy sftp://user@50.50.50.50/switchsoftware/efos-3.4.4.6.stk backup
Remote Password:*****

Mode..... SFTP
Set Server IP..... 50.50.50.50
Path..... /switchsoftware/
Filename..... efos-3.4.4.6.stk
Data Type..... Code
Destination Filename..... backup

Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y

File transfer in progress. Management access will be blocked for the
duration of the transfer. Please wait...
SFTP Code transfer starting...

File transfer operation completed successfully.

(IP_switch_A_1) #
```

2. Set the switch to boot from the backup partition on the next switch reboot:

```
boot system backup
```

```
(IP_switch_A_1) #boot system backup
Activating image backup ..

(IP_switch_A_1) #
```

3. Verify that the new boot image will be active on the next boot:

```
show bootvar
```

```
(IP_switch_A_1) #show bootvar
```

```
Image Descriptions
```

```
active :
```

```
backup :
```

```
Images currently available on Flash
```

unit	active	backup	current-active	next-active
1	3.4.4.2	3.4.4.6	3.4.4.2	3.4.4.6

```
(IP_switch_A_1) #
```

4. Save the configuration:

```
write memory
```

```
(IP_switch_A_1) #write memory
```

```
This operation may take a few minutes.
```

```
Management interfaces will not be available during this time.
```

```
Are you sure you want to save? (y/n) y
```

```
Configuration Saved!
```

```
(IP_switch_A_1) #
```

5. Reboot the switch:

```
reload
```

```
(IP_switch_A_1) #reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

6. Wait for the switch to reboot.



In rare scenarios the switch may fail to boot. Follow the [Steps to upgrade EFOS using the ONIE OS installation](#) to install the new image.

7. If you change the switch from EFOS 3.4.x.x to EFOS 3.7.x.x or vice versa then follow the following two procedures to apply the correct configuration (RCF):
 - a. [Resetting the Broadcom IP switch to factory defaults](#)
 - b. [Downloading and installing the Broadcom RCF files](#)
8. Repeat these steps on the remaining three IP switches in the MetroCluster IP configuration.

Steps to upgrade EFOS using the ONIE OS installation

You can perform the following steps if one EFOS version is FIPS compliant and the other EFOS version is non-FIPS compliant. These steps can be used to install the non-FIPS or FIPS compliant EFOS 3.7.x.x image from ONIE if the switch fails to boot.

Steps

1. Boot the switch into ONIE installation mode.

During boot, select ONIE when the following screen appears:

```
+-----+
| EFOS  |
| *ONIE |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
|       |
+-----+
```

After selecting "ONIE", the switch will then load and present you with the following choices:

```

+-----+
|*ONIE: Install OS
| ONIE: Rescue
| ONIE: Uninstall OS
| ONIE: Update ONIE
| ONIE: Embed ONIE
| DIAG: Diagnostic Mode
| DIAG: Burn-In Mode
|
|
|
|
|
+-----+

```

The switch now will boot into ONIE installation mode.

2. Stop the ONIE discovery and configure the ethernet interface

Once the following message appears press <enter> to invoke the ONIE console:

```

Please press Enter to activate this console. Info: eth0: Checking
link... up.
ONIE:/ #

```



The ONIE discovery will continue and messages will be printed to the console.

```

Stop the ONIE discovery
ONIE:/ # onie-discovery-stop
discover: installer mode detected.
Stopping: discover... done.
ONIE:/ #

```

3. Configure the ethernet interface and add the route using `ifconfig eth0 <ipAddress> netmask <netmask> up` and `route add default gw <gatewayAddress>`

```

ONIE:/ # ifconfig eth0 10.10.10.10 netmask 255.255.255.0 up
ONIE:/ # route add default gw 10.10.10.1

```

4. Verify that the server hosting the ONIE installation file is reachable:

```

ONIE:/ # ping 50.50.50.50
PING 50.50.50.50 (50.50.50.50): 56 data bytes
64 bytes from 50.50.50.50: seq=0 ttl=255 time=0.429 ms
64 bytes from 50.50.50.50: seq=1 ttl=255 time=0.595 ms
64 bytes from 50.50.50.50: seq=2 ttl=255 time=0.369 ms
^C
--- 50.50.50.50 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0.369/0.464/0.595 ms
ONIE:/ #

```

5. Install the new switch software

```

ONIE:/ # onie-nos-install http:// 50.50.50.50/Software/onie-installer-
x86_64
discover: installer mode detected.
Stopping: discover... done.
Info: Fetching http:// 50.50.50.50/Software/onie-installer-3.7.0.4 ...
Connecting to 50.50.50.50 (50.50.50.50:80)
installer          100% |*****| 48841k
0:00:00 ETA
ONIE: Executing installer: http:// 50.50.50.50/Software/onie-installer-
3.7.0.4
Verifying image checksum ... OK.
Preparing image archive ... OK.

```

The software will install and then reboot the switch. Let the switch reboot normally into the new EFOS version.

6. Verify that the new switch software is installed

show bootvar

```

(Routing) #show bootvar
Image Descriptions
active :
backup :
Images currently available on Flash
-----
unit      active      backup      current-active      next-active
-----
1      3.7.0.4      3.7.0.4      3.7.0.4      3.7.0.4
(Routing) #

```

7. Complete the installation

The switch will reboot with no configuration applied and reset to factory defaults. Follow the two procedures to configure the switch basic settings and apply the RCF file as outlined in the following two documents:

- a. Configure the switch basic settings. Follow step 4 and later: [Resetting the Broadcom IP switch to factory defaults](#)
- b. Create and apply the RCF file as outlined in [Downloading and installing the Broadcom RCF files](#)

Downloading and installing the Broadcom RCF files

You must generate and install the switch RCF file to each switch in the MetroCluster IP configuration.

Before you begin

This task requires file transfer software, such as FTP, TFTP, SFTP, or SCP, to copy the files to the switches.

About this task

These steps must be repeated on each of the IP switches in the MetroCluster IP configuration.

There are four RCF files, one for each of the four switches in the MetroCluster IP configuration. You must use the correct RCF files for the switch model you are using.

Switch	RCF file
IP_switch_A_1	v1.32_Switch-A1.txt
IP_switch_A_2	v1.32_Switch-A2.txt
IP_switch_B_1	v1.32_Switch-B1.txt
IP_switch_B_2	v1.32_Switch-B2.txt



The RCF files for EFOS version 3.4.4.6 or later 3.4.x.x. release and EFOS version 3.7.0.4 are different. You need to make sure that you have created the correct RCF files for the EFOS version that the switch is running.

EFOS version	RCF file version
3.4.x.x	v1.3x, v1.4x
3.7.x.x	v2.x

Steps

1. Generate the Broadcom RCF files for MetroCluster IP.
 - a. Download the [RcfFileGenerator for MetroCluster IP](#)
 - b. Generate the RCF file for your configuration using the RcfFileGenerator for MetroCluster IP.



Modifications to the RCF files after download are not supported.

2. Copy the RCF files to the switches:

a. Copy the RCF files to the first switch:

```
copy sftp://user@FTP-server-IP-address/RcfFiles/switch-specific-RCF/BES-53248_v1.32_Switch-A1.txt nvram:script BES-53248_v1.32_Switch-A1.scr
```

In this example, the "BES-53248_v1.32_Switch-A1.txt" RCF file is copied from the SFTP server at "50.50.50.50" to the local bootflash. You need to use the IP address of your TFTP/SFTP server and the file name of the RCF file that you need to install.

```
(IP_switch_A_1) #copy sftp://user@50.50.50.50/RcfFiles/BES-53248_v1.32_Switch-A1.txt nvram:script BES-53248_v1.32_Switch-A1.scr
```

```
Remote Password:*****
```

```
Mode..... SFTP
Set Server IP..... 50.50.50.50
Path..... /RcfFiles/
Filename..... BES-53248_v1.32_Switch-A1.txt
Data Type..... Config Script
Destination Filename..... BES-53248_v1.32_Switch-A1.scr
```

```
Management access will be blocked for the duration of the transfer
Are you sure you want to start? (y/n) y
```

```
File transfer in progress. Management access will be blocked for the
duration of the transfer. Please wait...
File transfer operation completed successfully.
```

```
Validating configuration script...
```

```
config
```

```
set clibanner
```

```
*****
*****
```

```
* NetApp Reference Configuration File (RCF)
```

```
*
```

```
* Switch : BES-53248
```

```
...
```

```
The downloaded RCF is validated. Some output is being logged here.
```

```
...
```

```
Configuration script validated.
```

```
File transfer operation completed successfully.
```

```
(IP_switch_A_1) #
```

b. Verify that the RCF file is saved as a script:

```
script list
```

```
(IP_switch_A_1) #script list

Configuration Script Name          Size(Bytes)  Date of Modification
-----
BES-53248_v1.32_Switch-A1.scr      852         2019 01 29 18:41:25

1 configuration script(s) found.
2046 Kbytes free.
(IP_switch_A_1) #
```

c. Apply the RCF script:

```
script apply BES-53248_v1.32_Switch-A1.scr
```

```
(IP_switch_A_1) #script apply BES-53248_v1.32_Switch-A1.scr

Are you sure you want to apply the configuration script? (y/n) y

config

set clibanner
"*****
*****

* NetApp Reference Configuration File (RCF)

*

* Switch      : BES-53248

...
The downloaded RCF is validated. Some output is being logged here.
...

Configuration script 'BES-53248_v1.32_Switch-A1.scr' applied.

(IP_switch_A_1) #
```

d. Save the configuration:

```
write memory
```

```
(IP_switch_A_1) #write memory
```

This operation may take a few minutes.
Management interfaces will not be available during this time.

```
Are you sure you want to save? (y/n) y
```

```
Configuration Saved!
```

```
(IP_switch_A_1) #
```

e. Reboot the switch:

```
reload
```

```
(IP_switch_A_1) #reload
```

```
Are you sure you would like to reset the system? (y/n) y
```

f. Repeat the previous steps for each of the other three switches, being sure to copy the matching RCF file to the corresponding switch.

3. Reload the switch:

```
reload
```

```
IP_switch_A_1# reload
```

4. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Disable unused ISL ports and port channels

NetApp recommends disabling unused ISL ports and port channels to avoid unnecessary health alerts.

1. Identify the unused ISL ports and port channels using the RCF file banner:



If the port is in breakout mode, the port name you specify in the command might be different than the name stated in the RCF banner. You can also use the RCF cabling files to find the port name.

For ISL port details

Run the command `show port all`.

For port channel details

Run the command `show port-channel all`.

2. Disable the unused ISL ports and port channels.

You must run the following commands for each identified unused port or port channel.

```
(SwtichA_1)> enable
(SwtichA_1)# configure
(SwtichA_1) (Config)# <port_name>
(SwtichA_1) (Interface 0/15)# shutdown
(SwtichA_1) (Interface 0/15)# end
(SwtichA_1)# write memory
```

Configure Cisco IP switches

Configure Cisco IP switches for cluster interconnect and backend MetroCluster IP connectivity

You must configure the Cisco IP switches for use as the cluster interconnect and for backend MetroCluster IP connectivity.

About this task

Several of the procedures in this section are independent procedures and you only need to execute those you are directed to or are relevant to your task.

Resetting the Cisco IP switch to factory defaults

Before installing any RCF file, you must erase the Cisco switch configuration and perform basic configuration. This procedure is required when you want to reinstall the same RCF file after a previous installation failed, or if you want to install a new version of an RCF file.

About this task

- You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.
- You must be connected to the switch using the serial console.
- This task resets the configuration of the management network.

Steps

1. Reset the switch to factory defaults:
 - a. Erase the existing configuration:

```
write erase
```

- b. Reload the switch software:

reload

The system reboots and enters the configuration wizard. During the boot, if you receive the prompt “Abort Auto Provisioning and continue with normal setup? (yes/no)[n]”, you should respond `yes` to proceed.

c. In the configuration wizard, enter the basic switch settings:

- Admin password
- Switch name
- Out-of-band management configuration
- Default gateway
- SSH service (RSA)

After completing the configuration wizard, the switch reboots.

d. When prompted, enter the user name and password to log in to the switch.

The following example shows the prompts and system responses when configuring the switch. The angle brackets (<<<) show where you enter the information.

```
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:y
**<<<**
```

```
Enter the password for "admin": password
Confirm the password for "admin": password
```

```
---- Basic System Configuration Dialog VDC: 1 ----
```

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

Please register Cisco Nexus3000 Family devices promptly with your supplier. Failure to register may affect response times for initial service calls. Nexus3000 devices must be registered to receive entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

You enter basic information in the next set of prompts, including the switch name, management address, and gateway, and select SSH with RSA.



This example shows the minimum information required to configure the RCF, additional options can be configured after the RCF has been applied. For example, you can configure SNMPv3, NTP, or SCP/SFTP after you have applied the RCF.

```

Would you like to enter the basic configuration dialog (yes/no): yes
Create another login account (yes/no) [n]:
Configure read-only SNMP community string (yes/no) [n]:
Configure read-write SNMP community string (yes/no) [n]:
Enter the switch name : switch-name **<<<
Continue with Out-of-band (mgmt0) management configuration?
(yes/no) [y]:
  Mgmt0 IPv4 address : management-IP-address **<<<
  Mgmt0 IPv4 netmask : management-IP-netmask **<<<
Configure the default gateway? (yes/no) [y]: y **<<<
  IPv4 address of the default gateway : gateway-IP-address **<<<
Configure advanced IP options? (yes/no) [n]:
Enable the telnet service? (yes/no) [n]:
Enable the ssh service? (yes/no) [y]: y **<<<
  Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
**<<<
  Number of rsa key bits <1024-2048> [1024]:
Configure the ntp server? (yes/no) [n]:
Configure default interface layer (L3/L2) [L2]:
Configure default switchport interface state (shut/noshut)
[noshut]: shut **<<<
  Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]:

```

The final set of prompts completes the configuration:

The following configuration will be applied:

```
password strength-check
switchname IP_switch_A_1
vrf context management
ip route 0.0.0.0/0 10.10.99.1
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address 10.10.99.10 255.255.255.0
no shutdown
```

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]:

```
2017 Jun 13 21:24:43 A1 %$ VDC-1 %$ %COPP-2-COPP_POLICY: Control-
Plane is protected with policy copp-system-p-policy-strict.
```

```
[#####] 100%
Copy complete.
```

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
.
.
.
IP_switch_A_1#
```

2. Save the configuration:

```
IP_switch-A-1# copy running-config startup-config
```

3. Reboot the switch and wait for the switch to reload:

```
IP_switch-A-1# reload
```

4. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Downloading and installing the Cisco switch NX-OS software

You must download the switch operating system file and RCF file to each switch in the MetroCluster IP configuration.

About this task

This task requires file transfer software, such as FTP, TFTP, SFTP, or SCP, to copy the files to the switches.

These steps must be repeated on each of the IP switches in the MetroCluster IP configuration.

You must use the supported switch software version.

[NetApp Hardware Universe](#)

Steps

1. Download the supported NX-OS software file.

[Cisco Software Download](#)

2. Copy the switch software to the switch:

```
copy sftp://root@server-ip-address/tftpboot/NX-OS-file-name bootflash: vrf
management
```

In this example, the nxos.7.0.3.I4.6.bin file and EPLD image is copied from SFTP server 10.10.99.99 to the local bootflash:

```

IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/nxos.7.0.3.I4.6.bin
bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/nxos.7.0.3.I4.6.bin
/bootflash/nxos.7.0.3.I4.6.bin
Fetching /tftpboot/nxos.7.0.3.I4.6.bin to /bootflash/nxos.7.0.3.I4.6.bin
/tftpboot/nxos.7.0.3.I4.6.bin          100% 666MB 7.2MB/s
01:32
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.

IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/n9000-
epld.9.3.5.img bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/n9000-epld.9.3.5.img /bootflash/n9000-
epld.9.3.5.img
Fetching /tftpboot/n9000-epld.9.3.5.img to /bootflash/n9000-
epld.9.3.5.img
/tftpboot/n9000-epld.9.3.5.img          161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.

```

3. Verify on each switch that the switch NX-OS files are present in each switch's bootflash directory:

```
dir bootflash:
```

The following example shows that the files are present on IP_switch_A_1:

```

IP_switch_A_1# dir bootflash:
      .
      .
      .
698629632   Jun 13 21:37:44 2017   nxos.7.0.3.I4.6.bin
      .
      .
      .

Usage for bootflash://sup-local
 1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#

```

4. Install the switch software:

install all nxos bootflash:nxos.version-number.bin

The switch will reload (reboot) automatically after the switch software has been installed.

The following example shows the software installation on IP_switch_A_1:

```

IP_switch_A_1# install all nxos bootflash:nxos.7.0.3.I4.6.bin
Installer will perform compatibility check first. Please wait.
Installer is forced disruptive

Verifying image bootflash:/nxos.7.0.3.I4.6.bin for boot variable "nxos".
[#####] 100% -- SUCCESS

Verifying image type.
[#####] 100% -- SUCCESS

Preparing "nxos" version info using image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS

Preparing "bios" version info using image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS           [#####] 100%
-- SUCCESS

Performing module support checks.           [#####] 100%
-- SUCCESS

Notifying services about system upgrade.     [#####] 100%

```

```
-- SUCCESS
```

```
Compatibility check is done:
```

Module	bootable	Impact	Install-type	Reason
1	yes	disruptive	reset	default upgrade is not hitless

```
Images will be upgraded according to following table:
```

Module Required	Image	Running-Version(pri:alt)	New-Version	Upg-
1	nxos	7.0(3)I4(1)	7.0(3)I4(6)	yes
1	bios	v04.24(04/21/2016)	v04.24(04/21/2016)	no

```
Switch will be reloaded for disruptive upgrade.
```

```
Do you want to continue with the installation (y/n)? [n] y
```

```
Install is in progress, please wait.
```

```
Performing runtime checks. [#####] 100% --  
SUCCESS
```

```
Setting boot variables.  
[#####] 100% -- SUCCESS
```

```
Performing configuration copy.  
[#####] 100% -- SUCCESS
```

```
Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.  
Warning: please do not remove or power off the module at this time.  
[#####] 100% -- SUCCESS
```

```
Finishing the upgrade, switch will reboot in 10 seconds.  
IP_switch_A_1#
```

5. Wait for the switch to reload and then log in to the switch.

After the switch has rebooted the login prompt is displayed:

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
.
MDP database restore in progress.
IP_switch_A_1#

The switch software is now installed.
```

6. Verify that the switch software has been installed:

```
show version
```

The following example shows the output:

```
IP_switch_A_1# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
.

Software
  BIOS: version 04.24
  NXOS: version 7.0(3)I4(6)   **<<< switch software version**
  BIOS compile time: 04/21/2016
  NXOS image file is: bootflash:///nxos.7.0.3.I4.6.bin
  NXOS compile time: 3/9/2017 22:00:00 [03/10/2017 07:05:18]

Hardware
  cisco Nexus 3132QV Chassis
  Intel(R) Core(TM) i3- CPU @ 2.50GHz with 16401416 kB of memory.
  Processor Board ID FOC20123GPS

  Device name: A1
  bootflash: 14900224 kB
  usb1: 0 kB (expansion flash)

Kernel uptime is 0 day(s), 0 hour(s), 1 minute(s), 49 second(s)

Last reset at 403451 usecs after Mon Jun 10 21:43:52 2017

Reason: Reset due to upgrade
System version: 7.0(3)I4(1)
Service:

plugin
  Core Plugin, Ethernet Plugin
IP_switch_A_1#
```

7. Upgrade the EPLD image and reboot the switch.

```

IP_switch_A_1# install epld bootflash:n9000-epld.9.3.5.img module 1
Compatibility check:
Module          Type          Upgradable    Impact        Reason
-----
1              SUP          Yes           disruptive    Module Upgradable

Retrieving EPLD versions.... Please wait.
Images will be upgraded according to following table:
Module  Type  EPLD          Running-Version  New-Version  Upg-
Required
-----
1  SUP  MI FPGA      0x07            0x07        No
1  SUP  IO FPGA      0x17            0x19        Yes
1  SUP  MI FPGA2     0x02            0x02        No

The above modules require upgrade.
The switch will be reloaded at the end of the upgrade
Do you want to continue (y/n) ? [n] y

Proceeding to upgrade Modules.

Starting Module 1 EPLD Upgrade

Module 1 : IO FPGA [Programming] : 100.00% (      64 of      64 sectors)
Module 1 EPLD upgrade is successful.
Module  Type  Upgrade-Result
-----
1  SUP  Success

EPLDs upgraded.

Module 1 EPLD upgrade is successful.

```

- After the switch reboot, log in again and verify that the new version of EPLD loaded successfully.

```
show version module 1 epld
```

- Repeat these steps on the remaining three IP switches in the MetroCluster IP configuration.

Downloading and installing the Cisco IP RCF files

You must generate and install the RCF file to each switch in the MetroCluster IP configuration.

About this task

This task requires file transfer software, such as FTP, TFTP, SFTP, or SCP, to copy the files to the switches.

These steps must be repeated on each of the IP switches in the MetroCluster IP configuration.

You must use the supported switch software version.

NetApp Hardware Universe

If you are using a QSFP-to-SFP+ adapter, you might need to configure the ISL port in native speed mode instead of breakout speed mode. Refer to your switch vendor documentation to determine the ISL port speed mode.

There are four RCF files, one for each of the four switches in the MetroCluster IP configuration. You must use the correct RCF files for the switch model you are using.

Switch	RCF file
IP_switch_A_1	NX3232_v1.80_Switch-A1.txt
IP_switch_A_2	NX3232_v1.80_Switch-A2.txt
IP_switch_B_1	NX3232_v1.80_Switch-B1.txt
IP_switch_B_2	NX3232_v1.80_Switch-B2.txt

Steps

1. Generate the Cisco RCF files for MetroCluster IP.
 - a. Download the [RcfFileGenerator for MetroCluster IP](#)
 - b. Generate the RCF file for your configuration using the RcfFileGenerator for MetroCluster IP.



Modifications to the RCF files after download are not supported.

2. Copy the RCF files to the switches:
 - a. Copy the RCF files to the first switch:

```
copy sftp://root@FTP-server-IP-address/tftpboot/switch-specific-RCF
bootflash: vrf management
```

In this example, the NX3232_v1.80_Switch-A1.txt RCF file is copied from the SFTP server at 10.10.99.99 to the local bootflash. You must use the IP address of your TFTP/SFTP server and the file name of the RCF file that you need to install.

```

IP_switch_A_1# copy
sftp://root@10.10.99.99/tftpboot/NX3232_v1.80_Switch-A1.txt
bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/NX3232_v1.80_Switch-A1.txt
/bootflash/NX3232_v1.80_Switch-A1.txt
Fetching /tftpboot/NX3232_v1.80_Switch-A1.txt to
/bootflash/NX3232_v1.80_Switch-A1.txt
/tftpboot/NX3232_v1.80_Switch-A1.txt          100% 5141      5.0KB/s
00:00
sftp> exit
Copy complete, now saving to disk (please wait)...
IP_switch_A_1#

```

b. Repeat the previous substep for each of the other three switches, being sure to copy the matching RCF file to the corresponding switch.

3. Verify on each switch that the RCF file is present in each switch's bootflash directory:

```
dir bootflash:
```

The following example shows that the files are present on IP_switch_A_1:

```

IP_switch_A_1# dir bootflash:
.
.
.
5514   Jun 13 22:09:05 2017  NX3232_v1.80_Switch-A1.txt
.
.
.

Usage for bootflash://sup-local
1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#

```

4. Configure the TCAM regions on Cisco 3132Q-V and Cisco 3232C switches.



Skip this step if you do not have Cisco 3132Q-V or Cisco 3232C switches.

a. On Cisco 3132Q-V switch, set the following TCAM regions:

```
conf t
hardware access-list tcam region span 0
hardware access-list tcam region racl 256
hardware access-list tcam region e-racl 256
hardware access-list tcam region qos 256
```

- b. On Cisco 3232C switch, set the following TCAM regions:

```
conf t
hardware access-list tcam region span 0
hardware access-list tcam region racl-lite 0
hardware access-list tcam region racl 256
hardware access-list tcam region e-racl 256
hardware access-list tcam region qos 256
```

- c. After setting the TCAM regions, save the configuration and reload the switch:

```
copy running-config startup-config
reload
```

5. Copy the matching RCF file from the local bootflash to the running configuration on each switch:

```
copy bootflash:switch-specific-RCF.txt running-config
```

6. Copy the RCF files from the running configuration to the startup configuration on each switch:

```
copy running-config startup-config
```

You should see output similar to the following:

```
IP_switch_A_1# copy bootflash:NX3232_v1.80_Switch-A1.txt running-config
IP_switch-A-1# copy running-config startup-config
```

7. Reload the switch:

```
reload
```

```
IP_switch_A_1# reload
```

8. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Setting Forward Error Correction for systems using 25-Gbps connectivity

If your system is configured using 25-Gbps connectivity, you need to set the Forward Error Correction (fec) parameter manually to off after applying the RCF file. The RCF file does not apply this setting.

About this task

The 25-Gbps ports must be cabled prior to performing this procedure.

[Platform port assignments for Cisco 3232C or Cisco 9336C switches](#)

This task only applies to platforms using 25-Gbps connectivity:

- AFF A300
- FAS 8200
- FAS 500f
- AFF A250

This task must be performed on all four switches in the MetroCluster IP configuration.

Steps

1. Set the fec parameter to off on each 25-Gbps port that is connected to a controller module, and then copy the running configuration to the startup configuration:
 - a. Enter configuration mode: `conf t`
 - b. Specify the 25-Gbps interface to configure: `interface interface-ID`
 - c. Set fec to off: `fec off`
 - d. Repeat the previous steps for each 25-Gbps port on the switch.
 - e. Exit configuration mode: `exit`

The following example shows the commands for interface Ethernet1/25/1 on switch IP_switch_A_1:

```
IP_switch_A_1# conf t
IP_switch_A_1(config)# interface Ethernet1/25/1
IP_switch_A_1(config-if)# fec off
IP_switch_A_1(config-if)# exit
IP_switch_A_1(config-if)# end
IP_switch_A_1# copy running-config startup-config
```

2. Repeat the previous step on the other three switches in the MetroCluster IP configuration.

Disable unused ISL ports and port channels

NetApp recommends disabling unused ISL ports and port channels to avoid unnecessary health alerts.

1. Identify the unused ISL ports and port channels:

```
show interface brief
```

2. Disable the unused ISL ports and port channels.

You must run the following commands for each identified unused port or port channel.

```
SwitchA_1# config t
Enter configuration commands, one per line. End with CNTL/Z.
SwitchA_1(config)# int Eth1/14
SwitchA_1(config-if)# shutdown
SwitchA_12(config-if)# exit
SwitchA_1(config-if)# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
```

Configure MACsec encryption on Cisco 9336C switches in a MetroCluster IP site

You can configure MACsec encryption on the WAN ISL ports that run between the sites. You must configure MACsec after applying the correct RCF file.



MACsec encryption can only be applied to the WAN ISL ports.

Configure MACsec encryption on Cisco 9336C switches

You must only configure MACsec encryption on the WAN ISL ports that run between the sites. You must configure MACsec after applying the correct RCF file.

Licensing requirements for MACsec

MACsec requires a security license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply for licenses, see the [Cisco NX-OS Licensing Guide](#)

Enable Cisco MACsec Encryption WAN ISLs in MetroCluster IP configurations

You can enable MACsec encryption for Cisco 9336C switches on the WAN ISLs in a MetroCluster IP configuration.

Steps

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Enable MACsec and MKA on the device:

```
feature macsec
```

```
IP_switch_A_1(config)# feature macsec
```

3. Copy the running configuration to the startup configuration:

```
copy running-config startup-config
```

```
IP_switch_A_1(config)# copy running-config startup-config
```

Configure a MACsec key chain and keys

You can create a MACsec key chain or keys on your configuration.

Key Lifetime and Hitless Key Rollover

A MACsec keychain can have multiple pre-shared keys (PSKs), each configured with a key ID and an optional lifetime. A key lifetime specifies at which time the key activates and expires. In the absence of a lifetime configuration, the default lifetime is unlimited. When a lifetime is configured, MKA rolls over to the next configured pre-shared key in the keychain after the lifetime is expired. The time zone of the key can be local or UTC. The default time zone is UTC. A key can roll over to a second key within the same keychain if you configure the second key (in the keychain) and configure a lifetime for the first key. When the lifetime of the first key expires, it automatically rolls over to the next key in the list. If the same key is configured on both sides of the link at the same time, then the key rollover is hitless (that is, the key rolls over without traffic interruption).

Steps

1. Enter the global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal  
IP_switch_A_1(config)#
```

2. To hide the encrypted key octet string, replace the string with a wildcard character in the output of the `show running-config` and `show startup-config` commands:

```
IP_switch_A_1(config)# key-chain macsec-psk no-show
```



The octet string is also hidden when you save the configuration to a file.

By default, PSK keys are displayed in encrypted format and can easily be decrypted. This command applies only to MACsec key chains.

3. Create a MACsec key chain to hold a set of MACsec keys and enter MACsec key chain configuration mode:

```
key chain name macsec
```

```
IP_switch_A_1(config)# key chain 1 macsec
IP_switch_A_1(config-macseckeychain)#
```

4. Create a MACsec key and enter MACsec key configuration mode:

```
key key-id
```

The range is from 1 to 32 hex digit key-string, and the maximum size is 64 characters.

```
IP_switch_A_1 switch(config-macseckeychain)# key 1000
IP_switch_A_1 (config-macseckeychain-macseckey)#
```

5. Configure the octet string for the key:

```
key-octet-string octet-string cryptographic-algorithm AES_128_CMAC |
AES_256_CMAC
```

```
IP_switch_A_1(config-macseckeychain-macseckey)# key-octet-string
abcdef0123456789abcdef0123456789abcdef0123456789abcdef0123456789
cryptographic-algorithm AES_256_CMAC
```



The octet-string argument can contain up to 64 hexadecimal characters. The octet key is encoded internally, so the key in clear text does not appear in the output of the `show running-config macsec` command.

6. Configure a send lifetime for the key (in seconds):

```
send-lifetime start-time duration duration
```

```
IP_switch_A_1(config-macseckeychain-macseckey)# send-lifetime 00:00:00
Oct 04 2020 duration 100000
```

By default, the device treats the start time as UTC. The start-time argument is the time of day and date that the key becomes active. The duration argument is the length of the lifetime in seconds. The maximum length is 2147483646 seconds (approximately 68 years).

7. Copy the running configuration to the startup configuration:

```
copy running-config startup-config
```

```
IP_switch_A_1(config)# copy running-config startup-config
```

8. Displays the keychain configuration:

show key chain name

```
IP_switch_A_1(config-macseckeychain-macseckey)# show key chain 1
```

Configure a MACsec policy

Steps

1. Enter global configuration mode:

configure terminal

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Create a MACsec policy:

macsec policy name

```
IP_switch_A_1(config)# macsec policy abc
IP_switch_A_1(config-macsec-policy)#
```

3. Configure one of the following ciphers, GCM-AES-128, GCM-AES-256, GCM-AES-XPN-128, or GCM-AES-XPN-256:

cipher-suite name

```
IP_switch_A_1(config-macsec-policy)# cipher-suite GCM-AES-256
```

4. Configure the key server priority to break the tie between peers during a key exchange:

key-server-priority number

```
switch(config-macsec-policy)# key-server-priority 0
```

5. Configure the security policy to define the handling of data and control packets:

security-policy security policy

Choose a security policy from the following options:

- must-secure — packets not carrying MACsec headers are dropped
- should-secure — packets not carrying MACsec headers are permitted (this is the default value)

```
IP_switch_A_1(config-macsec-policy)# security-policy should-secure
```

6. Configure the replay protection window so the secured interface does not accept a packet that is less than the configured window size: `window-size number`



The replay protection window size represents the maximum out-of-sequence frames that MACsec accepts and are not discarded. The range is from 0 to 596000000.

```
IP_switch_A_1(config-macsec-policy)# window-size 512
```

7. Configure the time in seconds to force an SAK rekey:

```
sak-expiry-time time
```

You can use this command to change the session key to a predictable time interval. The default is 0.

```
IP_switch_A_1(config-macsec-policy)# sak-expiry-time 100
```

8. Configure one of the following confidentiality offsets in the layer 2 frame where encryption begins:

```
conf-offset confidentiality offset
```

Choose from the following options:

- CONF-OFFSET-0.
- CONF-OFFSET-30.
- CONF-OFFSET-50.

```
IP_switch_A_1(config-macsec-policy)# conf-offset CONF-OFFSET-0
```



This command might be necessary for intermediate switches to use packet headers (dmac, smac, etype) like MPLS tags.

9. Copy the running configuration to the startup configuration:

```
copy running-config startup-config
```

```
IP_switch_A_1(config)# copy running-config startup-config
```

10. Display the MACsec policy configuration:

```
show macsec policy
```

```
IP_switch_A_1(config-macsec-policy)# show macsec policy
```

Enable Cisco MACsec encryption on the interfaces

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Select the interface that you configured with MACsec encryption.

You can specify the interface type and identity. For an Ethernet port, use ethernet slot/port.

```
IP_switch_A_1(config)# interface ethernet 1/15
switch(config-if)#
```

3. Add the keychain and policy to be configured on the interface to add the MACsec configuration:

```
macsec keychain keychain-name policy policy-name
```

```
IP_switch_A_1(config-if)# macsec keychain 1 policy abc
```

4. Repeat steps 1 and 2 on all interfaces where MACsec encryption is to be configured.
5. Copy the running configuration to the startup configuration:

```
copy running-config startup-config
```

```
IP_switch_A_1(config)# copy running-config startup-config
```

Disable Cisco MACsec Encryption WAN ISLs in MetroCluster IP configurations

You might need to disable MACsec encryption for Cisco 9336C switches on the WAN ISLs in a MetroCluster IP configuration.

Steps

1. Enter global configuration mode:

```
configure terminal
```

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Disable the MACsec configuration on the device:

```
macsec shutdown
```

```
IP_switch_A_1(config)# macsec shutdown
```



Selecting the “no” option restores the MACsec feature.

3. Select the interface that you already configured with MACsec.

You can specify the interface type and identity. For an Ethernet port, use ethernet slot/port.

```
IP_switch_A_1(config)# interface ethernet 1/15
switch(config-if)#
```

4. Remove the keychain and policy configured on the interface to remove the MACsec configuration:

```
no macsec keychain keychain-name policy policy-name
```

```
IP_switch_A_1(config-if)# no macsec keychain 1 policy abc
```

5. Repeat steps 3 and 4 on all interfaces where MACsec is configured.

6. Copy the running configuration to the startup configuration:

```
copy running-config startup-config
```

```
IP_switch_A_1(config)# copy running-config startup-config
```

Verifying the MACsec configuration

Steps

1. Repeat **all** of the previous procedures on the second switch within the configuration to establish a MACsec session.
2. Run the following commands to verify that both switches are successfully encrypted:
 - a. Run: `show macsec mka summary`
 - b. Run: `show macsec mka session`
 - c. Run: `show macsec mka statistics`

You can verify the MACsec configuration using the following commands:

Command	Displays information about...
<code>show macsec mka session interface typeslot/port number</code>	The MACsec MKA session for a specific interface or for all interfaces
<code>show key chain name</code>	The key chain configuration
<code>show macsec mka summary</code>	The MACsec MKA configuration
<code>show macsec policy policy-name</code>	The configuration for a specific MACsec policy or for all MACsec policies

Configure NVIDIA IP switches

Configure NVIDIA IP SN2100 switch for cluster interconnect and backend MetroCluster IP connectivity

You must configure the NVIDIA SN2100 IP switches for use as the cluster interconnect and for backend MetroCluster IP connectivity.

Reset the NVIDIA IP SN2100 switch to factory defaults

You can choose from the following methods to reset a switch to factory default settings.

- [Reset the switch using the RCF file option](#)
- [Download and install the Cumulus software](#)

Reset the switch using the RCF file option

Before installing a new RCF configuration you must revert the NVIDIA switch settings.

About this task

To restore the switch to default settings, run the RCF file with the `restoreDefaults` option. This option copies the original backed up files to their original location and then reboots the switch. After reboot, the switch comes online with the original configuration that existed when you first ran the RCF file to configure the switch.

The following configuration details are not reset:

- User and credential configuration
- Configuration of the management network port, eth0



All other configuration changes that occur during application of the RCF file are reverted to the original configuration.

Before you begin

- You must configure the switch according to [Download and install the NVIDIA RCF file](#). If you have not configured in this manner, or you have configured additional features before running the RCF file, you cannot use this procedure.

- You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.
- You must be connected to the switch with a serial console connection.
- This task resets the configuration of the management network.

Steps

1. Verify that the RCF configuration was successfully applied with the same or a compatible RCF file version and that the backup files exist.



The output can show backup files, preserved files, or both. If backup files or preserved files do not appear in the output, you cannot use this procedure.

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3
SN2100_v2.0.0_IP_switch_A_1.py
[sudo] password for cumulus:
>>> Opened RcfApplyLog
A RCF configuration has been successfully applied.
Backup files exist.
Preserved files exist.
Listing completion of the steps:
    Success: Step: 1: Performing Backup and Restore
    Success: Step: 2: updating MOTD file
    Success: Step: 3: Disabling apt-get
    Success: Step: 4: Disabling cdp
    Success: Step: 5: Adding lldp config
    Success: Step: 6: Creating interfaces
    Success: Step: 7: Configuring switch basic settings: Hostname,
SNMP
    Success: Step: 8: Configuring switch basic settings: bandwidth
allocation
    Success: Step: 9: Configuring switch basic settings: ecn
    Success: Step: 10: Configuring switch basic settings: cos and
dscp remark
    Success: Step: 11: Configuring switch basic settings: generic
egress cos mappings
    Success: Step: 12: Configuring switch basic settings: traffic
classification
    Success: Step: 13: Configuring LAG load balancing policies
    Success: Step: 14: Configuring the VLAN bridge
    Success: Step: 15: Configuring local cluster ISL ports
    Success: Step: 16: Configuring MetroCluster ISL ports
    Success: Step: 17: Configuring ports for MetroCluster-1, local
cluster and MetroCluster interfaces
    Success: Step: 18: Configuring ports for MetroCluster-2, local
cluster and MetroCluster interfaces
    Success: Step: 19: Configuring ports for MetroCluster-3, local
cluster and MetroCluster interfaces
    Success: Step: 20: Configuring L2FC for MetroCluster interfaces
    Success: Step: 21: Configuring the interface to UP
    Success: Step: 22: Final commit
    Success: Step: 23: Final reboot of the switch
Exiting ...
<<< Closing RcfApplyLog
cumulus@IP_switch_A_1:mgmt:~$

```

2. Run the RCF file with the option to restore defaults: `restoreDefaults`

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3
SN2100_v2.0.0_IP_switch_A_2.py restoreDefaults
[sudo] password for cumulus:
>>> Opened RcfApplyLog
Can restore from backup directory. Continuing.
This will reboot the switch !!!
Enter yes or no: yes

```

3. Respond 'yes' to the prompt. The switch reverts to the original configuration and reboots.
4. Wait for the switch to reboot.

The switch is reset and retains the initial configuration such as management network configuration and current credentials as they existed before applying the RCF file. After reboot, you can apply a new configuration by using the same or a different version of the RCF file.

Download and install the Cumulus software

About this task

Use these steps if you want to reset the switch completely by applying the Cumulus image.

Before you begin

- You must be connected to the switch with a serial console connection.
- The Cumulus switch software image is accessible through HTTP.



For more information about installing Cumulus Linux, see [Overview of installation and configuration for NVIDIA SN2100 switches](#)

- You must have the root password for `sudo` access to the commands.

Steps

1. From the Cumulus console download and queue the switch software installation with the command `onie-install -a -i` followed by the file path to the switch software:

In this example, the firmware file `cumulus-linux-4.4.3-mlx-amd64.bin` is copied from the HTTP server '50.50.50.50' to the local switch.

```

cumulus@IP_switch_A_1:mgmt:~$ sudo onie-install -a -i
http://50.50.50.50/switchsoftware/cumulus-linux-4.4.3-mlx-amd64.bin
Fetching installer: http://50.50.50.50/switchsoftware/cumulus-linux-
4.4.3-mlx-amd64.bin
Downloading URL: http://50.50.50.50/switchsoftware/cumulus-linux-4.4.3-
mlx-amd64.bin
#####
# 100.0%
Success: HTTP download complete.
tar: ./sysroot.tar: time stamp 2021-01-30 17:00:58 is 53895092.604407122

```

```
s in the future
tar: ./kernel: time stamp 2021-01-30 17:00:58 is 53895092.582826352 s in
the future
tar: ./initrd: time stamp 2021-01-30 17:00:58 is 53895092.509682557 s in
the future
tar: ./embedded-installer/bootloader/grub: time stamp 2020-12-10
15:25:16 is 49482950.509433937 s in the future
tar: ./embedded-installer/bootloader/init: time stamp 2020-12-10
15:25:16 is 49482950.509336507 s in the future
tar: ./embedded-installer/bootloader/uboot: time stamp 2020-12-10
15:25:16 is 49482950.509213637 s in the future
tar: ./embedded-installer/bootloader: time stamp 2020-12-10 15:25:16 is
49482950.509153787 s in the future
tar: ./embedded-installer/lib/init: time stamp 2020-12-10 15:25:16 is
49482950.509064547 s in the future
tar: ./embedded-installer/lib/logging: time stamp 2020-12-10 15:25:16 is
49482950.508997777 s in the future
tar: ./embedded-installer/lib/platform: time stamp 2020-12-10 15:25:16
is 49482950.508913317 s in the future
tar: ./embedded-installer/lib/utility: time stamp 2020-12-10 15:25:16 is
49482950.508847367 s in the future
tar: ./embedded-installer/lib/check-onie: time stamp 2020-12-10 15:25:16
is 49482950.508761477 s in the future
tar: ./embedded-installer/lib: time stamp 2020-12-10 15:25:47 is
49482981.508710647 s in the future
tar: ./embedded-installer/storage/blk: time stamp 2020-12-10 15:25:16 is
49482950.508631277 s in the future
tar: ./embedded-installer/storage/gpt: time stamp 2020-12-10 15:25:16 is
49482950.508523097 s in the future
tar: ./embedded-installer/storage/init: time stamp 2020-12-10 15:25:16
is 49482950.508437507 s in the future
tar: ./embedded-installer/storage/mbr: time stamp 2020-12-10 15:25:16 is
49482950.508371177 s in the future
tar: ./embedded-installer/storage/mtd: time stamp 2020-12-10 15:25:16 is
49482950.508293856 s in the future
tar: ./embedded-installer/storage: time stamp 2020-12-10 15:25:16 is
49482950.508243666 s in the future
tar: ./embedded-installer/platforms.db: time stamp 2020-12-10 15:25:16
is 49482950.508179456 s in the future
tar: ./embedded-installer/install: time stamp 2020-12-10 15:25:47 is
49482981.508094606 s in the future
tar: ./embedded-installer: time stamp 2020-12-10 15:25:47 is
49482981.508044066 s in the future
tar: ./control: time stamp 2021-01-30 17:00:58 is 53895092.507984316 s
in the future
tar: .: time stamp 2021-01-30 17:00:58 is 53895092.507920196 s in the
```

```
future
Staging installer image...done.
WARNING:
WARNING: Activating staged installer requested.
WARNING: This action will wipe out all system data.
WARNING: Make sure to back up your data.
WARNING:
Are you sure (y/N)? y
Activating staged installer...done.
Reboot required to take effect.
cumulus@IP_switch_A_1:mgmt:~$
```

2. Respond `y` to the prompt to confirm the installation when the image is downloaded and verified.
3. Reboot the switch to install the new software: `sudo reboot`

```
cumulus@IP_switch_A_1:mgmt:~$ sudo reboot
```



The switch reboots and enters the switch software installation which takes some time. When the installation is complete, the switch reboots and remains at the 'log-in' prompt.

4. Configure the basic switch settings
 - a. When the switch is booted and at the log-in prompt, log in and change the password.



The username is 'cumulus' and the default password is 'cumulus'.

```
Debian GNU/Linux 10 cumulus ttyS0

cumulus login: cumulus
Password:
You are required to change your password immediately (administrator
enforced)
Changing password for cumulus.
Current password:
New password:
Retype new password:
Linux cumulus 4.19.0-cl-1-amd64 #1 SMP Cumulus 4.19.206-1+cl4.4.3u1
(2021-12-18) x86_64

Welcome to NVIDIA Cumulus (R) Linux (R)

For support and online technical documentation, visit
http://www.cumulusnetworks.com/support

The registered trademark Linux (R) is used pursuant to a sublicense from
LMI,
the exclusive licensee of Linus Torvalds, owner of the mark on a world-
wide
basis.

cumulus@cumulus:~$
```

5. Configure the management network interface.

The commands you use depend on the switch firmware version you are running.



The following example commands configure the hostname as `IP_switch_A_1`, the IP address as `10.10.10.10`, the netmask as `255.255.255.0 (24)`, and the gateway address as `10.10.10.1`.

Cumulus 4.4.x

The following example commands configure the hostname, IP address, netmask, and gateway on a switch running Cumulus 4.4.x.

```
cumulus@cumulus:mgmt:~$ net add hostname IP_switch_A_1
cumulus@cumulus:mgmt:~$ net add interface eth0 ip address
10.0.10.10/24
cumulus@cumulus:mgmt:~$ net add interface eth0 ip gateway 10.10.10.1
cumulus@cumulus:mgmt:~$ net pending
```

```
.
.
.
```

```
cumulus@cumulus:mgmt:~$ net commit
```

```
.
.
.
```

```
net add/del commands since the last "net commit"
```

User Timestamp Command

```
cumulus 2021-05-17 22:21:57.437099 net add hostname Switch-A-1
cumulus 2021-05-17 22:21:57.538639 net add interface eth0 ip address
10.10.10.10/24
cumulus 2021-05-17 22:21:57.635729 net add interface eth0 ip gateway
10.10.10.1
```

```
cumulus@cumulus:mgmt:~$
```

Cumulus 5.4.x and later

The following example commands configure the hostname, IP address, netmask, and gateway on a switch running Cumulus 5.4.x. or later.

```
cumulus@cumulus:mgmt:~$ nv set system hostname IP_switch_A_1

cumulus@cumulus:mgmt:~$ nv set interface eth0 ip address
10.0.10.10/24

cumulus@cumulus:mgmt:~$ nv set interface eth0 ip gateway 10.10.10.1

cumulus@cumulus:mgmt:~$ nv config apply

cumulus@cumulus:mgmt:~$ nv config save
```

6. Reboot the switch using the `sudo reboot` command.

```
cumulus@cumulus:~$ sudo reboot
```

When the switch reboots, you can apply a new configuration using the steps in [Download and install the NVIDIA RCF file](#).

Download and install the NVIDIA RCF files

You must generate and install the switch RCF file to each switch in the MetroCluster IP configuration.

Before you begin

- You must have the root password for `sudo` access to the commands.
- The switch software is installed and the management network is configured.
- You followed the steps to initially install the switch by using either method 1 or method 2.
- You did not apply any additional configuration after the initial installation.



If you perform further configuration after resetting the switch and before applying the RCF file, you cannot use this procedure.

About this task

You must repeat these steps on each of the IP switches in the MetroCluster IP configuration (new installation) or on the replacement switch (switch replacement).

If you are using a QSFP-to-SFP+ adapter, you might need to configure the ISL port in native speed mode instead of breakout speed mode. Refer to your switch vendor documentation to determine the ISL port speed mode.

Steps

1. Generate the NVIDIA RCF files for MetroCluster IP.
 - a. Download the [RcfFileGenerator for MetroCluster IP](#).
 - b. Generate the RCF file for your configuration by using the RcfFileGenerator for MetroCluster IP.

c. Navigate to your home directory. If you are logged as 'cumulus', the file path is /home/cumulus.

```
cumulus@IP_switch_A_1:mgmt:~$ cd ~
cumulus@IP_switch_A_1:mgmt:~$ pwd
/home/cumulus
cumulus@IP_switch_A_1:mgmt:~$
```

d. Download the RCF file to this directory.

The following example shows that you use SCP to download the file

SN2100_v2.0.0_IP_switch_A_1.txt from server '50.50.50.50' to your home directory and save it as SN2100_v2.0.0_IP_switch_A_1.py:

```
cumulus@Switch-A-1:mgmt:~$ scp
username@50.50.50.50:/RcfFiles/SN2100_v2.0.0_IP_switch_A_1.txt
./SN2100_v2.0.0_IP_switch-A1.py
The authenticity of host '50.50.50.50 (50.50.50.50)' can't be
established.
RSA key fingerprint is
SHA256:B5gBtOmNZvdKiY+dPhh8=ZK9DaKG7g6sv+2gFlGVF8E.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '50.50.50.50' (RSA) to the list of known
hosts.
*****
**
Banner of the SCP server
*****
**
username@50.50.50.50's password:
SN2100_v2.0.0_IP_switch_A1.txt 100% 55KB 1.4MB/s 00:00
cumulus@IP_switch_A_1:mgmt:~$
```

2. Execute the RCF file.

The RCF file requires an option to apply one or more steps. Unless instructed by technical support, run the RCF file without the command line option. To verify the completion status of the various steps of the RCF file, use the option '-1' or 'all' to apply all (pending) steps.

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3
SN2100_v2.0.0_IP_switch_A_1.py
all
[sudo] password for cumulus:
The switch will be rebooted after the step(s) have been run.
Enter yes or no: yes

... the steps will apply - this is generating a lot of output ...

Running Step 24: Final reboot of the switch

... The switch will reboot if all steps applied successfully ...

```

3. If your configuration uses DAC cables, enable the DAC option on the switch ports:

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3 SN2100_v2.0.0-X10_Switch-
A1.py runCmd <switchport> DacOption [enable | disable]

```

The following example enables the DAC option for port swp7:

```

cumulus@IP_switch_A_1:mgmt:~$ sudo python3 SN2100_v2.00_Switch-A1.py
runCmd swp7 DacOption enable
Running cumulus version : 5.4.0
Running RCF file version : v2.00
Running command: Enabling the DacOption for port swp7
runCmd: 'nv set interface swp7 link fast-linkup on', ret: 0
runCmd: committed, ret: 0
Completion: SUCCESS
cumulus@IP_switch_A_1:mgmt:~$

```

4. Reboot the switch after enabling the DAC option on the switch ports:

```
sudo reboot
```



When you set the DAC option for multiple switch ports, you only need to reboot the switch once.

Set Forward Error Correction for systems using 25-Gbps connectivity

If your system is configured using 25-Gbps connectivity, set the Forward Error Correction (fec) parameter manually to off after applying the RCF. The RCF does not apply this setting.

About this task

- This task only applies to platforms using 25-Gbps connectivity. Refer to [Platform port assignments for NVIDIA supported SN2100 IP switches](#).
- This task must be performed on all four switches in the MetroCluster IP configuration.
- You must update each switch port individually, you cannot specify multiple ports or port ranges in the command.

Steps

1. Set the `fec` parameter to off for the first switch port that uses 25-Gbps connectivity:

```
sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport> fec off
```

2. Repeat the step for each 25-Gbps switch port that is connected to a controller module.

Set the switch port speed for the MetroCluster IP interfaces

About this task

- Use this procedure to set the switch port speed to 100G for the following systems:
 - AFF A70, AFF A90, AFF A1K, AFF C80
 - AFF A30, AFF C30, AFF A50, AFF C60
 - FAS50, FAS70, FAS90
- You must update each switch port individually, you cannot specify multiple ports or port ranges in the command.

Step

1. Use the RCF file with the `runCmd` option to set the speed. This applies the setting and saves the configuration.

The following commands set the speed for the MetroCluster interfaces `swp7` and `swp8`:

```
sudo python3 SN2100_v2.20 _Switch-A1.py runCmd swp7 speed 100
```

```
sudo python3 SN2100_v2.20 _Switch-A1.py runCmd swp8 speed 100
```

Example

```
cumulus@Switch-A-1:mgmt:~$ sudo python3 SN2100_v2.20_Switch-A1.py runCmd
swp7 speed 100
[sudo] password for cumulus: <password>
Running cumulus version : 5.4.0
Running RCF file version : v2.20
Running command: Setting switchport swp7 to 100G speed
runCmd: 'nv set interface swp7 link auto-negotiate off', ret: 0
runCmd: 'nv set interface swp7 link speed 100G', ret: 0
runCmd: committed, ret: 0
Completion: SUCCESS
cumulus@Switch-A-1:mgmt:~$
```

Disable unused ISL ports and port channels

NetApp recommends disabling unused ISL ports and port channels to avoid unnecessary health alerts. You must disable each port or port channel individually, you cannot specify multiple ports or port ranges in the command.

Steps

1. Identify the unused ISL ports and port channels using the RCF file banner:



If the port is in breakout mode, the port name you specify in the command might be different than the name stated in the RCF banner. You can also use the RCF cabling files to find the port name.

```
net show interface
```

2. Disable the unused ISL ports and port channels using the RCF file.

```

cumulus@mcc1-integrity-a1:mgmt:~$ sudo python3 SN2100_v2.0_IP_Switch-
A1.py runCmd
[sudo] password for cumulus:
    Running cumulus version   : 5.4.0
    Running RCF file version  : v2.0
Help for runCmd:
    To run a command execute the RCF script as follows:
    sudo python3 <script> runCmd <option-1> <option-2> <option-x>
    Depending on the command more or less options are required. Example
to 'up' port 'swp1'
    sudo python3 SN2100_v2.0_IP_Switch-A1.py runCmd swp1 up
Available commands:
    UP / DOWN the switchport
        sudo python3 SN2100_v2.0_IP_Switch-A1.py runCmd <switchport>
state <up | down>
    Set the switch port speed
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport>
speed <10 | 25 | 40 | 100 | AN>
    Set the fec mode on the switch port
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport>
fec <default | auto | rs | baser | off>
    Set the [localISL | remoteISL] to 'UP' or 'DOWN' state
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd [localISL |
remoteISL] state [up | down]
    Set the option on the port to support DAC cables. This option
does not support port ranges.
    You must reload the switch after changing this option for
the required ports. This will disrupt traffic.
    This setting requires Cumulus 5.4 or a later 5.x release.
        sudo python3 SN2100_v2.0_Switch-A1.py runCmd <switchport>
DacOption [enable | disable]
cumulus@mcc1-integrity-a1:mgmt:~$

```

The following example command disables port "swp14":

```
sudo python3 SN2100_v2.0_Switch-A1.py runCmd swp14 state down
```

Repeat this step for each identified unused port or port channel.

Install the Ethernet Switch Health Monitor configuration file for a NVIDIA SN2100 MetroCluster IP switch

To configure Ethernet switch health monitoring on NVIDIA Ethernet switches, follow this procedure.

These instructions apply if NVIDIA X190006-PE and X190006-PI switches are not detected properly, which can be confirmed by running `system switch ethernet show` and checking if **OTHER** is shown for your model.

To identify your NVIDIA switch model, find the part-number with the command `nv show platform hardware` for NVIDIA CL 5.8 and earlier or `nv show platform` for later versions.



These steps are also recommended if you want health monitoring and log collection to work as intended when using NVIDIA CL 5.11.x with the following ONTAP releases. While health monitoring and log collection might still function without these steps, following them ensures everything operates correctly.

- 9.10.1P20, 9.11.1P18, 9.12.1P16, 9.13.1P8, 9.14.1, 9.15.1 and later patch releases

Before you begin

- Make sure that the ONTAP cluster is up and running.
- Enable SSH on the switch to use all of the features available in CSHM.
- Clear the `/mroot/etc/cshm_nod/nod_sign/` directory on all nodes:

- a. Enter the nodeshell:

```
system node run -node <name>
```

- b. Change to advanced privilege:

```
priv set advanced
```

- c. List the configuration files in the `/etc/cshm_nod/nod_sign` directory. If the directory exists and contains configuration files, it lists the file names.

```
ls /etc/cshm_nod/nod_sign
```

- d. Delete all configuration files corresponding to your connected switch models.

If you are unsure, remove all configuration files for the supported models listed above, then download and install the latest configuration files for those same models.

```
rm /etc/cshm_nod/nod_sign/<filename>
```

- e. Confirm that the deleted configuration files are no longer in the directory:

```
ls /etc/cshm_nod/nod_sign
```

Steps

1. Download the Ethernet switch health monitor configuration zip file based on the corresponding ONTAP release version. This file is available from the [NVIDIA Ethernet switches](#) page.
 - a. On the NVIDIA SN2100 Software download page, select **Nvidia CSHM File**.
 - b. On the Caution/Must read page, select the check box to agree.
 - c. On the End User License Agreement page, select the check box to agree and click **Accept & Continue**.
 - d. On the Nvidia CSHM File - Download page, select the applicable configuration file. The following files are available:

ONTAP 9.15.1 and later

- MSN2100-CB2FC-v1.4.zip
- MSN2100-CB2RC-v1.4.zip
- X190006-PE-v1.4.zip
- X190006-PI-v1.4.zip

ONTAP 9.11.1 through 9.14.1

- MSN2100-CB2FC_PRIOR_R9.15.1-v1.4.zip
- MSN2100-CB2RC_PRIOR_R9.15.1-v1.4.zip
- X190006-PE_PRIOR_9.15.1-v1.4.zip
- X190006-PI_PRIOR_9.15.1-v1.4.zip

2. Upload the applicable zip file to your internal web server.
3. Access the advanced mode setting from one of the ONTAP systems in the cluster.

```
set -privilege advanced
```

4. Run the switch health monitor configure command.

```
cluster1::> system switch ethernet configure-health-monitor
```

5. Verify that the command output ends with the following text for your ONTAP version:

ONTAP 9.15.1 and later

Ethernet switch health monitoring installed the configuration file.

ONTAP 9.11.1 through 9.14.1

SHM installed the configuration file.

ONTAP 9.10.1

CSHM downloaded package processed successfully.

If an error occurs, contact NetApp support.

6. Wait up to twice the Ethernet switch health monitor polling interval, found by running `system switch ethernet polling-interval show`, before completing the next step.
7. Run the command `system switch ethernet configure-health-monitor show` on the ONTAP system and make sure that the cluster switches are discovered with the monitored field set to **True** and the serial number field not showing **Unknown**.

```
cluster1::> system switch ethernet configure-health-monitor show
```



If your model is still showing **OTHER** after applying the configuration file, contact NetApp support.

See the [system switch ethernet configure-health-monitor](#) command for further details.

What's next?

[Configure switch health monitoring.](#)

Monitor MetroCluster IP switch health

Learn about switch health monitoring in a MetroCluster IP configuration

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes.

Important notes for configuring CSHM in a MetroCluster IP configuration

This section contains the generic steps for configuring SNMPv3 and log collection on Cisco, Broadcom, and NVIDIA SN2100 switches. You must follow the steps for a switch firmware version that is supported in a MetroCluster IP configuration. Refer to the [Hardware Universe](#) to verify the supported firmware versions.

In a MetroCluster configuration, you configure health monitoring on the local cluster switches only.

For log collection with Broadcom and Cisco switches, a new user should be created on the switch for each cluster with log collection enabled. In a MetroCluster configuration, this means that MetroCluster 1, MetroCluster 2, MetroCluster 3, and MetroCluster 4 all require a separate user to be configured on the switches. These switches do not support multiple SSH keys for the same user. Any additional log collection setup performed overwrites any pre-existing SSH keys for the user.

Before you configure the CSHM, you should disable unused ISLs to avoid any unnecessary ISL alerts.

Configure SNMPv3 to monitor the health of MetroCluster IP switches

In MetroCluster IP configurations, you can configure SNMPv3 to monitor the health of IP switches.

This procedure shows the generic steps for configuring SNMPv3 on a switch. Some of the switch firmware versions listed might not be supported in a MetroCluster IP configuration.

You must follow the steps for a switch firmware version that is supported in a MetroCluster IP configuration. Refer to the [Hardware Universe](#) to verify the supported firmware versions.

- SNMPv3 is only supported on ONTAP 9.12.1 and later.
- ONTAP 9.13.1P12, 9.14.1P9, 9.15.1P5, 9.16.1 and later versions fix these two issues:
 - [For ONTAP health monitoring of Cisco switches, SNMPv2 traffic might still be seen after switching to SNMPv3 for monitoring](#)
 - [False-positive switch fan and power alerts when SNMP failures occur](#)



About this task

The following commands are used to configure an SNMPv3 username on **Broadcom**, **Cisco**, and **NVIDIA** switches:

Broadcom switches

Configure an SNMPv3 username NETWORK-OPERATOR on Broadcom BES-53248 switches.

- For **no authentication**:

```
snmp-server user SNMPv3UserNoAuth NETWORK-OPERATOR noauth
```

- For **MD5/SHA authentication**:

```
snmp-server user SNMPv3UserAuth NETWORK-OPERATOR [auth-md5|auth-sha]
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
snmp-server user SNMPv3UserAuthEncrypt NETWORK-OPERATOR [auth-  
md5|auth-sha] [priv-aes128|priv-des]
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp  
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
cluster1::*> system switch ethernet modify -device DEVICE -snmp-version  
SNMPv3 -community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp status
```

```
(sw1) (Config)# snmp-server user <username> network-admin auth-md5
<password> priv-aes128 <password>
```

```
(cs1) (Config)# show snmp user snmp
```

Name	Group Name	Auth Meth	Priv Meth	Remote Engine ID
<username>	network-admin	MD5	AES128	8000113d03d8c497710bee

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name <username>
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance

Device Name: sw1
IP Address: 10.228.136.24
SNMP Version: SNMPv2c
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: cshml!
Model Number: BES-53248
Switch Network: cluster-network
Software Version: 3.9.0.2
Reason For Not Monitoring: None <---- should
display this if SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored ?: true
Serial Number of the Device: QTFCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>

```

4. After waiting the CSHM polling period, verify that the serial number is populated for the Ethernet switch.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
      Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance
      Device Name: sw1
      IP Address: 10.228.136.24
      SNMP Version: SNMPv3
      Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: <username>
      Model Number: BES-53248
      Switch Network: cluster-network
      Software Version: 3.9.0.2
      Reason For Not Monitoring: None <---- should
display this if SNMP settings are valid
      Source Of Switch Version: CDP/ISDP
      Is Monitored?: true
      Serial Number of the Device: QTFCU3826001C
      RCF Version: v1.8X2 for

Cluster/HA/RDMA

```

Cisco switches

Configure an SNMPv3 username SNMPv3_USER on Cisco 9336C-FX2 switches:

- For **no authentication**:

```
snmp-server user SNMPv3_USER NoAuth
```

- For **MD5/SHA authentication**:

```
snmp-server user SNMPv3_USER auth [md5|sha] AUTH-PASSWORD
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
snmp-server user SNMPv3_USER AuthEncrypt auth [md5|sha] AUTH-
PASSWORD priv aes-128 PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp user
```

```
(sw1) (Config)# snmp-server user SNMPv3User auth md5 <auth_password>
priv aes-128 <priv_password>
```

```
(sw1) (Config)# show snmp user
```

```
-----
-----
```

SNMP USERS

```
-----
-----
```

User	Auth	Priv(enforce)	Groups
acl_filter			
admin	md5	des(no)	network-admin
SNMPv3User	md5	aes-128(no)	network-operator

```
-----
-----
```

NOTIFICATION TARGET USERS (configured for sending V3 Inform)

```
-----
-----
```

User	Auth	Priv
------	------	------

```
(sw1) (Config)#
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -is-monitoring-enabled-admin true
```

```
cluster1::*> security login create -user-or-group-name <username>
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1" -instance

                Device Name: sw1
                IP Address: 10.231.80.212
                SNMP Version: SNMPv2c
                Is Discovered: true
                SNMPv2c Community String or SNMPv3 Username: cshml!
                Model Number: N9K-C9336C-FX2
                Switch Network: cluster-network
                Software Version: Cisco Nexus
Operating System (NX-OS) Software, Version 9.3(7)
                Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
                Source Of Switch Version: CDP/ISDP
                Is Monitored?: true
                Serial Number of the Device: QTFCU3826001C
                RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>
cluster1::*>

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```

system switch ethernet polling-interval show

```

```

cluster1::*> system switch ethernet polling-interval show
                Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance

                Device Name: sw1
                IP Address: 10.231.80.212
                SNMP Version: SNMPv3
                Is Discovered: true
                SNMPv2c Community String or SNMPv3 Username: SNMPv3User
                Model Number: N9K-C9336C-FX2
                Switch Network: cluster-network
                Software Version: Cisco Nexus
                Operating System (NX-OS) Software, Version 9.3(7)
                Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
                Source Of Switch Version: CDP/ISDP
                Is Monitored?: true
                Serial Number of the Device: QTFCU3826001C
                RCF Version: v1.8X2 for
                Cluster/HA/RDMA

cluster1::*>

```

NVIDIA - CL 5.4.0

Configure an SNMPv3 username SNMPv3_USER on NVIDIA SN2100 switches running CLI 5.4.0:

- For **no authentication**:

```
nv set service snmp-server username SNMPv3_USER auth-none
```

- For **MD5/SHA authentication**:

```
nv set service snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD
```

- For **MD5/SHA authentication with AES/DES encryption**:

```
nv set service snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
net show snmp status
```

```
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          4318
Version 1 and 2c Community String Configured
Version 3 Usernames     Not Configured
-----

cumulus@sw1:~$
cumulus@sw1:~$ net add snmp-server username SNMPv3User auth-md5
<password> encrypt-aes <password>
cumulus@sw1:~$ net commit
--- /etc/snmp/snmpd.conf      2020-08-02 21:09:34.686949282 +0000
+++ /run/nclu/snmp/snmpd.conf 2020-08-11 00:13:51.826126655 +0000
@@ -1,26 +1,28 @@
# Auto-generated config file: do not edit. #
agentaddress udp:@mgmt:161
agentxperms 777 777 snmp snmp
agentxsocket /var/agentx/master
createuser _snmptrapusernameX
+createuser SNMPv3User MD5 <password> AES <password>
ifmib_max_num_ifaces 500
iquerysecname _snmptrapusernameX
master agentx
monitor -r 60 -o laNames -o laErrorMessage "laTable" laErrorFlag != 0
pass -p 10 1.3.6.1.2.1.1.1 /usr/share/snmp/sysDescr_pass.py
pass_persist 1.2.840.10006.300.43
```

```

/usr/share/snmp/ieee8023_lag_pp.py
  pass_persist 1.3.6.1.2.1.17 /usr/share/snmp/bridge_pp.py
  pass_persist 1.3.6.1.2.1.31.1.1.1.18
/usr/share/snmp/snmpifAlias_pp.py
  pass_persist 1.3.6.1.2.1.47 /usr/share/snmp/entity_pp.py
  pass_persist 1.3.6.1.2.1.99 /usr/share/snmp/entity_sensor_pp.py
  pass_persist 1.3.6.1.4.1.40310.1 /usr/share/snmp/resq_pp.py
  pass_persist 1.3.6.1.4.1.40310.2
/usr/share/snmp/cl_drop_cntrs_pp.py
  pass_persist 1.3.6.1.4.1.40310.3 /usr/share/snmp/cl_poe_pp.py
  pass_persist 1.3.6.1.4.1.40310.4 /usr/share/snmp/bgpun_pp.py
  pass_persist 1.3.6.1.4.1.40310.5 /usr/share/snmp/cumulus-status.py
  pass_persist 1.3.6.1.4.1.40310.6 /usr/share/snmp/cumulus-sensor.py
  pass_persist 1.3.6.1.4.1.40310.7 /usr/share/snmp/vrf_bgpun_pp.py
+rocommunity cshml! default
  rouser _snmptrapusernameX
+rouser SNMPv3User priv
  sysobjectid 1.3.6.1.4.1.40310
  sysservices 72
-rocommunity cshml! default

```

net add/del commands since the last "net commit"

```

User          Timestamp          Command
-----
-----
SNMPv3User   2020-08-11 00:13:51.826987 net add snmp-server username
SNMPv3User auth-md5 <password> encrypt-aes <password>

```

```

cumulus@sw1:~$
cumulus@sw1:~$ net show snmp status
Simple Network Management Protocol (SNMP) Daemon.
-----
Current Status          active (running)
Reload Status           enabled
Listening IP Addresses  all vrf mgmt
Main snmpd PID          24253
Version 1 and 2c Community String Configured
Version 3 Usernames     Configured      <---- Configured
here
-----
cumulus@sw1:~$

```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name SNMPv3User -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name SNMPv3User
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)"
-instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
                                     Device Name: sw1
(b8:59:9f:09:7c:22)
                                     IP Address: 10.231.80.212
                                     SNMP Version: SNMPv2c
                                     Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: cshml!
                                     Model Number: MSN2100-CB2FC
                                     Switch Network: cluster-network
                                     Software Version: Cumulus Linux
version 5.4.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored ?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
      RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username
SNMPv3User

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: SNMPv3User
Model Number: MSN2100-CB2FC
Switch Network: cluster-network
Software Version: Cumulus Linux
version 5.4.0 running on Mellanox Technologies Ltd. MSN2100
Reason For Not Monitoring: None
Source Of Switch Version: LLDP
Is Monitored?: true
Serial Number of the Device: MT2110X06399 <----
serial number to check
RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

```

NVIDIA - CL 5.11.0

Configure an SNMPv3 username SNMPv3_USER on NVIDIA SN2100 switches running CLI 5.11.0:

- For no authentication:

```
nv set system snmp-server username SNMPv3_USER auth-none
```

- For MD5/SHA authentication:

```
nv set system snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD
```

- For MD5/SHA authentication with AES/DES encryption:

```
nv set system snmp-server username SNMPv3_USER [auth-md5|auth-sha]
AUTH-PASSWORD [encrypt-aes|encrypt-des] PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
security login create -user-or-group-name SNMPv3_USER -application snmp
-authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
system switch ethernet modify -device DEVICE -snmp-version SNMPv3
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
nv show system snmp-server
```

```
cumulus@sw1:~$ nv show system snmp-server
                                applied
-----
[username]                       SNMPv3_USER
[username]                       limiteduser1
[username]                       testuserauth
[username]                       testuserauthaes
[username]                       testusernoauth
trap-link-up
  check-frequency                 60
trap-link-down
  check-frequency                 60
[listening-address]              all
[readonly-community]             $nvsec$94d69b56e921aec1790844eb53e772bf
state                             enabled
cumulus@sw1:~$
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name SNMPv3User -application
snmp -authentication-method usm -remote-switch-ipaddress
10.231.80.212
```

```
cluster1::*> security login create -user-or-group-name SNMPv3User
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212
```

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha, sha2-256)

[none]: **md5**

Enter the authentication protocol password (minimum 8 characters long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)

[none]: **aes128**

Enter privacy protocol password (minimum 8 characters long):

Enter privacy protocol password again:

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1 (b8:59:9f:09:7c:22)"
-instance
```

```

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
                                     Device Name: sw1
(b8:59:9f:09:7c:22)
                                     IP Address: 10.231.80.212
                                     SNMP Version: SNMPv2c
                                     Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
      Community String or SNMPv3 Username: cshml!
      Model Number: MSN2100-CB2FC
      Switch Network: cluster-network
      Software Version: Cumulus Linux
version 5.11.0 running on Mellanox Technologies Ltd. MSN2100
      Reason For Not Monitoring: None
      Source Of Switch Version: LLDP
      Is Monitored ?: true
      Serial Number of the Device: MT2110X06399 <----
serial number to check
      RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -snmp-version SNMPv3 -community-or-username
SNMPv3User

```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

```

cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1
(b8:59:9f:09:7c:22)" -instance
Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
DEPRECATED-Community String or SNMPv3 Username: -
Community String or SNMPv3 Username: SNMPv3User
Model Number: MSN2100-CB2FC
Switch Network: cluster-network
Software Version: Cumulus Linux
version 5.11.0 running on Mellanox Technologies Ltd. MSN2100
Reason For Not Monitoring: None
Source Of Switch Version: LLDP
Is Monitored?: true
Serial Number of the Device: MT2110X06399 <----
serial number to check
RCF Version: MSN2100-RCF-v1.9X6-
Cluster-LLDP Aug-18-2022

```

Configure log collection on a MetroCluster IP switch

In a MetroCluster IP configuration, you can configure log collection to collect switch logs for debugging purposes.



On Broadcom and Cisco switches, a new user is required for each cluster with log collection. For example, MetroCluster 1, MetroCluster 2, MetroCluster 3, and MetroCluster 4 all require a separate user to be configured on the switches. Multiple SSH keys for the same user is not supported.

About this task

The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes. This procedure guides you through the process of setting up collection, requesting detailed **Support** logs, and enabling an hourly collection of **Periodic** data that is collected by AutoSupport.

NOTE: If you enable FIPS mode, you must complete the following:



1. Regenerate SSH keys on the switch using the vendor instructions.
2. Regenerate SSH keys in ONTAP using `debug system regenerate-systemshell-key-pair`
3. Re-run log collection setup routine using the `system switch ethernet log setup-password` command

Before you begin

- The user must have access to the switch `show` commands. If these are not available, create a new user and grant the necessary permissions to the user.
- Switch health monitoring must be enabled for the switch. Verify this by ensuring the `Is Monitored:` field is set to **true** in the output of the `system switch ethernet show` command.
- For log collection with Broadcom and Cisco switches:
 - The local user must have network admin privileges.
 - A new user should be created on the switch for each cluster setup with log collection enabled. These switches do not support multiple SSH keys for the same user. Any additional log collection setup performed overwrites any pre-existing SSH keys for the user.
- For support log collection with NVIDIA switches, the *user* for log collection must be permitted to run the `cl-support` command without having to provide a password. To allow this usage, run the command:

```
echo '<user> ALL = NOPASSWD: /usr/cumulus/bin/cl-support' | sudo EDITOR='tee -a' visudo -f /etc/sudoers.d/cumulus
```

Steps

ONTAP 9.15.1 and later

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

NOTE: If answering **y** to the user specification prompt, make sure that the user has the necessary permissions as outlined in [Before you begin](#).

```
system switch ethernet log setup-password
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1
```

```
cs2
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```



For CL 5.11.1, create the user **cumulus** and respond **y** to the following prompt: Would you like to specify a user other than admin for log collection? {y|n}: **y**

2. Enable periodic log collection:

```
system switch ethernet log modify -device <switch-name> -periodic  
-enabled true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -periodic
-enabled true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

cs1: Periodic log collection has been scheduled to run every hour.

```
cluster1::*> system switch ethernet log modify -device cs2 -periodic
-enabled true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

cs2: Periodic log collection has been scheduled to run every hour.

```
cluster1::*> system switch ethernet log show
```

	Periodic	Periodic
Support		
Switch	Log Enabled	Log State
Log State		
cs1	true	scheduled
never-run		
cs2	true	scheduled
never-run		

2 entries were displayed.

3. Request support log collection:

```
system switch ethernet log collect-support-log -device <switch-name>
```

```
cluster1::*> system switch ethernet log collect-support-log -device
cs1
```

```
cs1: Waiting for the next Ethernet switch polling cycle to begin
support collection.
```

```
cluster1::*> system switch ethernet log collect-support-log -device
cs2
```

```
cs2: Waiting for the next Ethernet switch polling cycle to begin
support collection.
```

```
cluster1::*> *system switch ethernet log show
```

	Periodic	Periodic
Support		
Switch	Log Enabled	Log State
Log State		

cs1	false	halted
initiated		

cs2	true	scheduled
initiated		

```
2 entries were displayed.
```

4. To view all details of log collection, including the enablement, status message, previous timestamp and filename of periodic collection, the request status, status message, and previous timestamp and filename of support collection, use the following:

```
system switch ethernet log show -instance
```

```

cluster1::*> system switch ethernet log show -instance

                Switch Name: cs1
    Periodic Log Enabled: true
        Periodic Log Status: Periodic log collection has been
scheduled to run every hour.
    Last Periodic Log Timestamp: 3/11/2024 11:02:59
        Periodic Log Filename: cluster1:/mroot/etc/log/shm-
cluster-info.tgz
    Support Log Requested: false
        Support Log Status: Successfully gathered support logs
- see filename for their location.
    Last Support Log Timestamp: 3/11/2024 11:14:20
        Support Log Filename: cluster1:/mroot/etc/log/shm-
cluster-log.tgz

                Switch Name: cs2
    Periodic Log Enabled: false
        Periodic Log Status: Periodic collection has been
halted.
    Last Periodic Log Timestamp: 3/11/2024 11:05:18
        Periodic Log Filename: cluster1:/mroot/etc/log/shm-
cluster-info.tgz
    Support Log Requested: false
        Support Log Status: Successfully gathered support logs
- see filename for their location.
    Last Support Log Timestamp: 3/11/2024 11:18:54
        Support Log Filename: cluster1:/mroot/etc/log/shm-
cluster-log.tgz
2 entries were displayed.

```

ONTAP 9.14.1 and earlier

1. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

NOTE: If answering `y` to the user specification prompt, make sure that the user has the necessary permissions as outlined in [Before you begin](#).

```
system switch ethernet log setup-password
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1
```

```
cs2
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2
```

```
Would you like to specify a user other than admin for log  
collection? {y|n}: n
```

```
Enter the password: <enter switch password>
```

```
Enter the password again: <enter switch password>
```



For CL 5.11.1, create the user **cumulus** and respond **y** to the following prompt: Would you like to specify a user other than admin for log collection? {y|n}: **y**

2. To request support log collection and enable periodic collection, run the following command. This starts both types of log collection: the detailed Support logs and an hourly collection of Periodic data.

```
system switch ethernet log modify -device <switch-name> -log-request  
true
```

```
cluster1::*> system switch ethernet log modify -device cs1 -log
-request true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

```
cluster1::*> system switch ethernet log modify -device cs2 -log
-request true
```

Do you want to modify the cluster switch log collection configuration? {y|n}: [n] **y**

Enabling cluster switch log collection.

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```



If any error statuses are reported by the log collection feature (visible in the output of `system switch ethernet log show`), see [Troubleshoot log collection](#) for further details.

Manage the monitoring of Ethernet switches in a MetroCluster IP configuration

In most cases, Ethernet switches are automatically discovered by ONTAP and monitored by CSHM. The Reference Configuration File (RCF) applied to the switch, among other things, enables the Cisco Discovery Protocol (CDP) and/or the Link Layer Discovery Protocol (LLDP). However, you might need to manually add a switch that is not discovered or remove a switch that is no longer in use. You can also stop active monitoring while retaining the switch in the configuration, such as during maintenance.

Create a switch entry so that ONTAP can monitor it

About this task

Use the `system switch ethernet create` command to manually configure and enable monitoring for a specified Ethernet switch. This is helpful if ONTAP does not add the switch automatically, or if you previously removed the switch and want to re-add it.

```
system switch ethernet create -device DeviceName -address 1.2.3.4 -snmp
-version SNMPv2c -community-or-username cshml! -model NX3132V -type
cluster-network
```

A typical example is adding a switch named [DeviceName], with IP address 1.2.3.4, and SNMPv2c credentials set to **cs hm1!**. Use `-type storage-network` instead of `-type cluster-network` if you are configuring a storage switch.

Disable monitoring without deleting the switch

If you want to pause or stop monitoring for a certain switch, but still retain it for future monitoring, modify its `is-monitoring-enabled-admin` parameter instead of deleting it.

For example:

```
system switch ethernet modify -device DeviceName -is-monitoring-enabled
-admin false
```

This lets you preserve switch details and configuration without generating new alerts or re-discoveries.

Remove a switch you no longer need

Use `system switch ethernet delete` to delete a switch that has been disconnected or is no longer required:

```
system switch ethernet delete -device DeviceName
```

By default, this command succeeds only if ONTAP does not currently detect the switch through CDP or LLDP. To remove a discovered switch, use the `-force` parameter:

```
system switch ethernet delete -device DeviceName -force
```

When `-force` is used, the switch might be re-added automatically if ONTAP detects it again.

Verify Ethernet switch monitoring in a MetroCluster IP configuration

The Ethernet switch health monitor (CSHM) automatically attempts to monitor the switches that it discovers; however, monitoring might not happen automatically if the switches are not configured correctly. You should verify that the health monitor is properly configured to monitor your switches.

Confirm monitoring of the connected Ethernet switches

About this task

To confirm that the connected Ethernet switches are being monitored, run:

```
system switch ethernet show
```

If the `Model` column displays **OTHER** or the `IS Monitored` field displays **false**, then ONTAP cannot monitor the switch. A value of **OTHER** typically indicates that ONTAP does not support that switch for health

monitoring.

The `IS Monitored` field is set to **false** for the reason specified in the `Reason` field.



If a switch is not listed in the command output, ONTAP likely has not discovered it. Confirm that the switch is cabled correctly. If necessary, you can add the switch manually. See [Manage the monitoring of Ethernet Switches](#) for further details.

Confirm firmware and RCF versions are up to date

Make sure that the switch is running the latest supported firmware and that a compatible reference configuration file (RCF) has been applied. More information is available on the [NetApp Support Downloads page](#).

By default, the health monitor uses SNMPv2c with the community string **csbm1!** for monitoring, but SNMPv3 can also be configured.

If you need to change the default SNMPv2c community string, make sure that the desired SNMPv2c community string has been configured on the switch.

```
system switch ethernet modify -device SwitchA -snmp-version SNMPv2c  
-community-or-username newCommunity!
```



See [Optional: Configure SNMPv3](#) for details on configuring SNMPv3 for use.

Confirm management network connection

Verify that the switch's management port is connected to the management network.

A correct management port connection is required for ONTAP to perform SNMP queries and log collection.

Related information

- [Troubleshoot alerts](#)

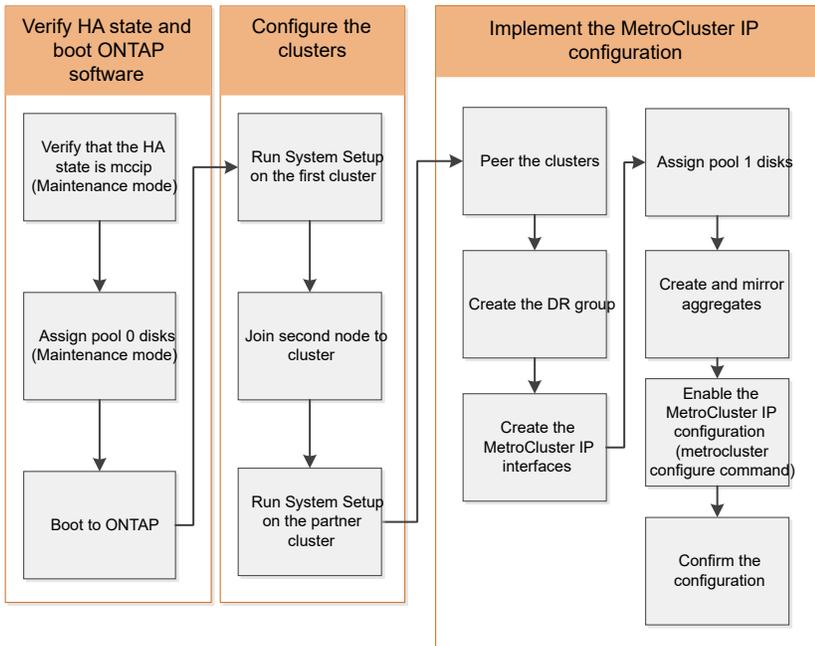
Configure the MetroCluster software in ONTAP

Configure the MetroCluster software using the CLI

Set up the ONTAP nodes and clusters in the MetroCluster IP configuration

You must set up each node in the MetroCluster configuration in ONTAP, including the node-level configurations and the configuration of the nodes into two sites. You must also implement the MetroCluster relationship between the two sites.

If a controller module fails during configuration, refer to [Controller module failure scenarios during MetroCluster installation](#).



Configure eight-node MetroCluster IP configurations

An eight-node MetroCluster configuration consists of two DR groups. To configure the first DR group, complete the tasks in this section. After you have configured the first DR group, you can follow the steps to [expand a four-node MetroCluster IP configuration to an eight-node configuration](#).

Gather the required information for your MetroCluster IP configuration

You need to gather the required IP addresses for the controller modules before you begin the configuration process.

You can use these links to download csv files and fill in the tables with your site-specific information.

[MetroCluster IP setup worksheet, site_A](#)

[MetroCluster IP setup worksheet, site_B](#)

Compare ONTAP standard cluster and MetroCluster configurations

The configuration of the nodes in each cluster in a MetroCluster configuration is similar to that of nodes in a standard cluster.

The MetroCluster configuration is built on two standard clusters. Physically, the configuration must be symmetrical, with each node having the same hardware configuration, and all of the MetroCluster components must be cabled and configured. However, the basic software configuration for nodes in a MetroCluster configuration is the same as that for nodes in a standard cluster.

Configuration step	Standard cluster configuration	MetroCluster configuration
Configure management, cluster, and data LIFs on each node.	Same in both types of clusters	

Configure the root aggregate.	Same in both types of clusters	
Set up the cluster on one node in the cluster.	Same in both types of clusters	
Join the other node to the cluster.	Same in both types of clusters	
Create a mirrored root aggregate.	Optional	Required
Peer the clusters.	Optional	Required
Enable the MetroCluster configuration.	Does not apply	Required

Verify the HA configuration state of your controller and chassis components in a MetroCluster IP configuration

In a MetroCluster IP configuration, you must verify that the ha-config state of the controller and chassis components is set to “mccip” so that they boot up properly. Although this value should be preconfigured on systems received from the factory, you should still verify the setting before proceeding.

If the HA state of the controller module and chassis is incorrect, you cannot configure the MetroCluster without re-initializing the node. You must correct the setting using this procedure, and then initialize the system by using one of the following procedures:



- In a MetroCluster IP configuration, follow the steps in [Restore system defaults on a controller module](#).
- In a MetroCluster FC configuration, follow the steps in [Restore system defaults and configuring the HBA type on a controller module](#).

Before you begin

Verify that the system is in Maintenance mode.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The correct HA state depends on your MetroCluster configuration.

MetroCluster configuration type	HA state for all components...
Eight or four node MetroCluster FC configuration	mcc
Two-node MetroCluster FC configuration	mcc-2n

Eight or four node MetroCluster IP configuration	mccip
--	-------

- If the displayed system state of the controller is not correct, set the correct HA state for your configuration on the controller module:

MetroCluster configuration type	Command
Eight or four node MetroCluster FC configuration	ha-config modify controller mcc
Two-node MetroCluster FC configuration	ha-config modify controller mcc-2n
Eight or four node MetroCluster IP configuration	ha-config modify controller mccip

- If the displayed system state of the chassis is not correct, set the correct HA state for your configuration on the chassis:

MetroCluster configuration type	Command
Eight or four node MetroCluster FC configuration	ha-config modify chassis mcc
Two-node MetroCluster FC configuration	ha-config modify chassis mcc-2n
Eight or four node MetroCluster IP configuration	ha-config modify chassis mccip

- Boot the node to ONTAP:

```
boot_ontap
```

- Repeat this entire procedure to verify the HA state on each node in the MetroCluster configuration.

Restore system defaults on a controller module before setting up a MetroCluster IP configuration

Reset and restore defaults on the controller modules.

- At the LOADER prompt, return environmental variables to their default setting: `set-defaults`
- Boot the node to the boot menu: `boot_ontap menu`

After you run this command, wait until the boot menu is shown.

- Clear the node configuration:

- If you are using systems configured for ADP, select option 9a from the boot menu, and respond `no` when prompted.



This process is disruptive.

The following screen shows the boot menu prompt:

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.
- (10) Set Onboard Key Manager recovery secrets.
- (11) Configure node for external key management.

Selection (1-11)? 9a

...

```
##### WARNING: AGGREGATES WILL BE DESTROYED #####  
This is a disruptive operation that applies to all the disks  
that are attached and visible to this node.
```

Before proceeding further, make sure that:

The aggregates visible from this node do not contain data that needs to be preserved.

This option (9a) has been executed or will be executed on the HA partner node (and DR/DR-AUX partner nodes if applicable), prior to reinitializing any system in the HA-pair or MetroCluster configuration.

The HA partner node (and DR/DR-AUX partner nodes if applicable) is currently waiting at the boot menu.

Do you want to abort this operation (yes/no)? no

- If your system is not configured for ADP, type `wipeconfig` at the boot menu prompt, and then press Enter.

The following screen shows the boot menu prompt:

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.

Selection (1-9)? wipeconfig

This option deletes critical system configuration, including cluster membership.

Warning: do not run this option on a HA node that has been taken over.

Are you sure you want to continue?: yes

Rebooting to finish wipeconfig request.

Manually assign drives to pool 0 in a MetroCluster IP configuration

If you did not receive the systems pre-configured from the factory, you might have to manually assign the pool 0 drives. Depending on the platform model and whether the system is using ADP, you must manually assign drives to pool 0 for each node in the MetroCluster IP configuration. The procedure you use depends on the version of ONTAP you are using.

Manually assigning drives for pool 0 (ONTAP 9.4 and later)

If the system has not been pre-configured at the factory and does not meet the requirements for automatic drive assignment, you must manually assign the pool 0 drives.

About this task

This procedure applies to configurations running ONTAP 9.4 or later.

To determine if your system requires manual disk assignment, you should review [Considerations for automatic drive assignment and ADP systems in ONTAP 9.4 and later](#).

You perform these steps in Maintenance mode. The procedure must be performed on each node in the configuration.

Examples in this section are based on the following assumptions:

- node_A_1 and node_A_2 own drives on:
 - site_A-shelf_1 (local)
 - site_B-shelf_2 (remote)

- node_B_1 and node_B_2 own drives on:
 - site_B-shelf_1 (local)
 - site_A-shelf_2 (remote)

Steps

1. Display the boot menu:

```
boot_ontap menu
```

2. Select Option 9a and respond `no` when prompted.

The following screen shows the boot menu prompt:

```
Please choose one of the following:
```

```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 9a
```

```
...
```

```
##### WARNING: AGGREGATES WILL BE DESTROYED #####
This is a disruptive operation that applies to all the disks
that are attached and visible to this node.
```

```
Before proceeding further, make sure that:
```

```
The aggregates visible from this node do not contain
data that needs to be preserved.
```

```
This option (9a) has been executed or will be executed
on the HA partner node (and DR/DR-AUX partner nodes if
applicable), prior to reinitializing any system in the
HA-pair or MetroCluster configuration.
```

```
The HA partner node (and DR/DR-AUX partner nodes if
applicable) is currently waiting at the boot menu.
Do you want to abort this operation (yes/no)? no
```

3. When the node restarts, press Ctrl-C when prompted to display the boot menu and then select the option for **Maintenance mode boot**.
4. In Maintenance mode, manually assign drives for the local aggregates on the node:

```
disk assign disk-id -p 0 -s local-node-sysid
```

The drives should be assigned symmetrically, so each node has an equal number of drives. The following steps are for a configuration with two storage shelves at each site.

- a. When configuring node_A_1, manually assign drives from slot 0 to 11 to pool0 of node A1 from site_A-shelf_1.
 - b. When configuring node_A_2, manually assign drives from slot 12 to 23 to pool0 of node A2 from site_A-shelf_1.
 - c. When configuring node_B_1, manually assign drives from slot 0 to 11 to pool0 of node B1 from site_B-shelf_1.
 - d. When configuring node_B_2, manually assign drives from slot 12 to 23 to pool0 of node B2 from site_B-shelf_1.
5. Exit Maintenance mode:

```
halt
```

6. Display the boot menu:

```
boot_ontap menu
```

7. Repeat these steps on the other nodes in the MetroCluster IP configuration.
8. Select Option **4** from the boot menu on both nodes and let the system boot.
9. Proceed to [Setting up ONTAP](#).

Manually assigning drives for pool 0 (ONTAP 9.3)

If you have at least two disk shelves for each node, you use ONTAP's auto-assignment functionality to automatically assign the local (pool 0) disks.

About this task

While the node is in Maintenance mode, you must first assign a single disk on the appropriate shelves to pool 0. ONTAP then automatically assigns the rest of the disks on the shelf to the same pool. This task is not required on systems received from the factory, which have pool 0 to contain the pre-configured root aggregate.

This procedure applies to configurations running ONTAP 9.3.

This procedure is not required if you received your MetroCluster configuration from the factory. Nodes from the factory are configured with pool 0 disks and root aggregates.

This procedure can be used only if you have at least two disk shelves for each node, which allows shelf-level autoassignment of disks. If you cannot use shelf-level autoassignment, you must manually assign your local disks so that each node has a local pool of disks (pool 0).

These steps must be performed in Maintenance mode.

Examples in this section assume the following disk shelves:

- node_A_1 owns disks on:
 - site_A-shelf_1 (local)
 - site_B-shelf_2 (remote)
- node_A_2 is connected to:
 - site_A-shelf_3 (local)
 - site_B-shelf_4 (remote)
- node_B_1 is connected to:
 - site_B-shelf_1 (local)
 - site_A-shelf_2 (remote)
- node_B_2 is connected to:
 - site_B-shelf_3 (local)
 - site_A-shelf_4 (remote)

Steps

1. Manually assign a single disk for root aggregate on each node:

```
disk assign disk-id -p 0 -s local-node-sysid
```

The manual assignment of these disks allows the ONTAP autoassignment feature to assign the rest of the disks on each shelf.

- a. On node_A_1, manually assign one disk from local site_A-shelf_1 to pool 0.
 - b. On node_A_2, manually assign one disk from local site_A-shelf_3 to pool 0.
 - c. On node_B_1, manually assign one disk from local site_B-shelf_1 to pool 0.
 - d. On node_B_2, manually assign one disk from local site_B-shelf_3 to pool 0.
2. Boot each node at site A, using option 4 on the boot menu:

You should complete this step on a node before proceeding to the next node.

- a. Exit Maintenance mode:

```
halt
```

- b. Display the boot menu:

```
boot_ontap menu
```

- c. Select option 4 from the boot menu and proceed.

3. Boot each node at site B, using option 4 on the boot menu:

You should complete this step on a node before proceeding to the next node.

- a. Exit Maintenance mode:

```
halt
```

- b. Display the boot menu:

`boot_ontap` menu

- c. Select option 4 from the boot menu and proceed.

Set up ONTAP nodes in a MetroCluster IP configuration

After you boot each node, you are prompted to perform basic node and cluster configuration. After configuring the cluster, you return to the ONTAP CLI to create aggregates and create the MetroCluster configuration.

Before you begin

- You must have cabled the MetroCluster configuration.

If you need to netboot the new controllers, see [Netboot the new controller modules](#).

About this task

This task must be performed on both clusters in the MetroCluster configuration.

Steps

1. Power up each node at the local site if you have not already done so and let them all boot completely.

If the system is in Maintenance mode, you need to issue the `halt` command to exit Maintenance mode, and then issue the `boot_ontap` command to boot the system and get to cluster setup.

2. On the first node in each cluster, proceed through the prompts to configure the cluster.
 - a. Enable the AutoSupport tool by following the directions provided by the system.

The output should be similar to the following:

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".

To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp Technical

Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and

resolution should a problem occur on your system.

For further information on AutoSupport, see:

<http://support.netapp.com/autosupport/>

Type yes to confirm and continue {yes}: yes

.
.
.

b. Configure the node management interface by responding to the prompts.

The prompts are similar to the following:

```
Enter the node management interface port [e0M]:
Enter the node management interface IP address: 172.17.8.229
Enter the node management interface netmask: 255.255.254.0
Enter the node management interface default gateway: 172.17.8.1
A node management interface on port e0M with IP address 172.17.8.229
has been created.
```

c. Create the cluster by responding to the prompts.

The prompts are similar to the following:

```
Do you want to create a new cluster or join an existing cluster?
{create, join}:
create
```

```
Do you intend for this node to be used as a single node cluster?
{yes, no} [no]:
no
```

```
Existing cluster interface configuration found:
```

```
Port MTU IP Netmask
e0a 1500 169.254.18.124 255.255.0.0
e1a 1500 169.254.184.44 255.255.0.0
```

```
Do you want to use this configuration? {yes, no} [yes]: no
```

```
System Defaults:
```

```
Private cluster network ports [e0a,e1a].
Cluster port MTU values will be set to 9000.
Cluster interface IP addresses will be automatically generated.
```

```
Do you want to use these defaults? {yes, no} [yes]: no
```

```
Enter the cluster administrator's (username "admin") password:
```

```
Retype the password:
```

```
Step 1 of 5: Create a Cluster
```

```
You can type "back", "exit", or "help" at any question.
```

```
List the private cluster network ports [e0a,e1a]:
```

```
Enter the cluster ports' MTU size [9000]:
```

```
Enter the cluster network netmask [255.255.0.0]: 255.255.254.0
```

```
Enter the cluster interface IP address for port e0a: 172.17.10.228
```

```
Enter the cluster interface IP address for port e1a: 172.17.10.229
```

```
Enter the cluster name: cluster_A
```

```
Creating cluster cluster_A
```

```
Starting cluster support services ...
```

```
Cluster cluster_A has been created.
```

- d. Add licenses, set up a Cluster Administration SVM, and enter DNS information by responding to the prompts.

The prompts are similar to the following:

```
Step 2 of 5: Add Feature License Keys
You can type "back", "exit", or "help" at any question.

Enter an additional license key []:

Step 3 of 5: Set Up a Vserver for Cluster Administration
You can type "back", "exit", or "help" at any question.

Enter the cluster management interface port [e3a]:
Enter the cluster management interface IP address: 172.17.12.153
Enter the cluster management interface netmask: 255.255.252.0
Enter the cluster management interface default gateway: 172.17.12.1

A cluster management interface on port e3a with IP address
172.17.12.153 has been created. You can use this address to connect
to and manage the cluster.

Enter the DNS domain names: lab.netapp.com
Enter the name server IP addresses: 172.19.2.30
DNS lookup for the admin Vserver will use the lab.netapp.com domain.

Step 4 of 5: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.

SFO will be enabled when the partner joins the cluster.

Step 5 of 5: Set Up the Node
You can type "back", "exit", or "help" at any question.

Where is the controller located []: svl
```

- e. Enable storage failover and set up the node by responding to the prompts.

The prompts are similar to the following:

```
Step 4 of 5: Configure Storage Failover (SFO)
You can type "back", "exit", or "help" at any question.
```

```
SFO will be enabled when the partner joins the cluster.
```

```
Step 5 of 5: Set Up the Node
You can type "back", "exit", or "help" at any question.
```

```
Where is the controller located []: site_A
```

- f. Complete the configuration of the node, but do not create data aggregates.

You can use ONTAP System Manager, pointing your web browser to the cluster management IP address (<https://172.17.12.153>).

[Cluster management using System Manager \(ONTAP 9.7 and earlier\)](#)

[ONTAP System Manager \(Version 9.7 and later\)](#)

- g. Configure the Service Processor (SP):

[Configure the SP/BMC network](#)

[Use a Service Processor with System Manager - ONTAP 9.7 and earlier](#)

3. Boot the next controller and join it to the cluster, following the prompts.
4. Confirm that nodes are configured in high-availability mode:

```
storage failover show -fields mode
```

If not, you must configure HA mode on each node, and then reboot the nodes:

```
storage failover modify -mode ha -node localhost
```



The expected configuration state of HA and storage failover is as follows:

- HA mode is configured but storage failover is not enabled.
- HA takeover capability is disabled.
- HA interfaces are offline.
- HA mode, storage failover, and interfaces are configured later in the process.

5. Confirm that you have four ports configured as cluster interconnects:

```
network port show
```

The MetroCluster IP interfaces are not configured at this time and do not appear in the command output.

The following example shows two cluster ports on node_A_1:

```
cluster_A::*> network port show -role cluster

Node: node_A_1

Ignore

Health
Speed(Mbps) Health

Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status

-----
-----

e4a      Cluster      Cluster      up    9000  auto/40000  healthy
false

e4e      Cluster      Cluster      up    9000  auto/40000  healthy
false

Node: node_A_2

Ignore

Health
Speed(Mbps) Health

Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status

-----
-----

e4a      Cluster      Cluster      up    9000  auto/40000  healthy
false

e4e      Cluster      Cluster      up    9000  auto/40000  healthy
```

```
false
```

```
4 entries were displayed.
```

6. Repeat these steps on the partner cluster.

What to do next

Return to the ONTAP command-line interface and complete the MetroCluster configuration by performing the tasks that follow.

Configure ONTAP clusters in a MetroCluster IP configuration

You must peer the clusters, mirror the root aggregates, create a mirrored data aggregate, and then issue the command to implement the MetroCluster operations.

About this task

Before you run `metrocluster configure`, HA mode and DR mirroring are not enabled and you might see an error message related to this expected behavior. You enable HA mode and DR mirroring later when you run the command `metrocluster configure` to implement the configuration.

Disabling automatic drive assignment (if doing manual assignment in ONTAP 9.4)

In ONTAP 9.4, if your MetroCluster IP configuration has fewer than four external storage shelves per site, you must disable automatic drive assignment on all nodes and manually assign drives.

About this task

This task is not required in ONTAP 9.5 and later.

This task does not apply to an AFF A800 system with an internal shelf and no external shelves.

[Considerations for automatic drive assignment and ADP systems in ONTAP 9.4 and later](#)

Steps

1. Disable automatic drive assignment:

```
storage disk option modify -node <node_name> -autoassign off
```

2. You need to issue this command on all nodes in the MetroCluster IP configuration.

Verifying drive assignment of pool 0 drives

You must verify that the remote drives are visible to the nodes and have been assigned correctly.

About this task

Automatic assignment depends on the storage system platform model and drive shelf arrangement.

[Considerations for automatic drive assignment and ADP systems in ONTAP 9.4 and later](#)

Steps

1. Verify that pool 0 drives are assigned automatically:

disk show

The following example shows the "cluster_A" output for an AFF A800 system with no external shelves.

One quarter (8 drives) were automatically assigned to "node_A_1" and one quarter were automatically assigned to "node_A_2". The remaining drives will be remote (pool 1) drives for "node_B_1" and "node_B_2".

```
cluster_A::*> disk show
Disk Owner Usable Size Disk Shelf Bay Container Type Container Name
-----
node_A_1:0n.12 1.75TB 0 12 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.13 1.75TB 0 13 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.14 1.75TB 0 14 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.15 1.75TB 0 15 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.16 1.75TB 0 16 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.17 1.75TB 0 17 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.18 1.75TB 0 18 SSD-NVM shared aggr0
node_A_1
node_A_1:0n.19 1.75TB 0 19 SSD-NVM shared -
node_A_1
node_A_2:0n.0 1.75TB 0 0 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.1 1.75TB 0 1 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.2 1.75TB 0 2 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.3 1.75TB 0 3 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.4 1.75TB 0 4 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.5 1.75TB 0 5 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.6 1.75TB 0 6 SSD-NVM shared
aggr0_node_A_2_0 node_A_2
node_A_2:0n.7 1.75TB 0 7 SSD-NVM shared -
node_A_2
node_A_2:0n.24 - 0 24 SSD-NVM unassigned - -
```

```

node_A_2:0n.25 - 0 25 SSD-NVM unassigned - -
node_A_2:0n.26 - 0 26 SSD-NVM unassigned - -
node_A_2:0n.27 - 0 27 SSD-NVM unassigned - -
node_A_2:0n.28 - 0 28 SSD-NVM unassigned - -
node_A_2:0n.29 - 0 29 SSD-NVM unassigned - -
node_A_2:0n.30 - 0 30 SSD-NVM unassigned - -
node_A_2:0n.31 - 0 31 SSD-NVM unassigned - -
node_A_2:0n.36 - 0 36 SSD-NVM unassigned - -
node_A_2:0n.37 - 0 37 SSD-NVM unassigned - -
node_A_2:0n.38 - 0 38 SSD-NVM unassigned - -
node_A_2:0n.39 - 0 39 SSD-NVM unassigned - -
node_A_2:0n.40 - 0 40 SSD-NVM unassigned - -
node_A_2:0n.41 - 0 41 SSD-NVM unassigned - -
node_A_2:0n.42 - 0 42 SSD-NVM unassigned - -
node_A_2:0n.43 - 0 43 SSD-NVM unassigned - -
32 entries were displayed.

```

The following example shows the "cluster_B" output:

```

cluster_B::> disk show
          Usable      Disk          Container  Container
Disk      Size      Shelf Bay Type      Type      Name
Owner
-----
-----

Info: This cluster has partitioned disks. To get a complete list of
spare disk
capacity use "storage aggregate show-spare-disks".
node_B_1:0n.12  1.75TB    0    12  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.13  1.75TB    0    13  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.14  1.75TB    0    14  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.15  1.75TB    0    15  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.16  1.75TB    0    16  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.17  1.75TB    0    17  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.18  1.75TB    0    18  SSD-NVM shared  aggr0
node_B_1
node_B_1:0n.19  1.75TB    0    19  SSD-NVM shared  -
node_B_1

```

```

node_B_2:0n.0      1.75TB      0      0      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.1      1.75TB      0      1      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.2      1.75TB      0      2      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.3      1.75TB      0      3      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.4      1.75TB      0      4      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.5      1.75TB      0      5      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.6      1.75TB      0      6      SSD-NVM shared
aggr0_node_B_1_0 node_B_2
node_B_2:0n.7      1.75TB      0      7      SSD-NVM shared      -
node_B_2
node_B_2:0n.24     -            0      24     SSD-NVM unassigned  -      -
node_B_2:0n.25     -            0      25     SSD-NVM unassigned  -      -
node_B_2:0n.26     -            0      26     SSD-NVM unassigned  -      -
node_B_2:0n.27     -            0      27     SSD-NVM unassigned  -      -
node_B_2:0n.28     -            0      28     SSD-NVM unassigned  -      -
node_B_2:0n.29     -            0      29     SSD-NVM unassigned  -      -
node_B_2:0n.30     -            0      30     SSD-NVM unassigned  -      -
node_B_2:0n.31     -            0      31     SSD-NVM unassigned  -      -
node_B_2:0n.36     -            0      36     SSD-NVM unassigned  -      -
node_B_2:0n.37     -            0      37     SSD-NVM unassigned  -      -
node_B_2:0n.38     -            0      38     SSD-NVM unassigned  -      -
node_B_2:0n.39     -            0      39     SSD-NVM unassigned  -      -
node_B_2:0n.40     -            0      40     SSD-NVM unassigned  -      -
node_B_2:0n.41     -            0      41     SSD-NVM unassigned  -      -
node_B_2:0n.42     -            0      42     SSD-NVM unassigned  -      -
node_B_2:0n.43     -            0      43     SSD-NVM unassigned  -      -
32 entries were displayed.

cluster_B::>

```

Peering the clusters

The clusters in the MetroCluster configuration must be in a peer relationship so that they can communicate with each other and perform the data mirroring essential to MetroCluster disaster recovery.

Related information

[Cluster and SVM peering express configuration](#)

[Considerations when using dedicated ports](#)

[Considerations when sharing data ports](#)

Configuring intercluster LIFs for cluster peering

You must create intercluster LIFs on ports used for communication between the MetroCluster partner clusters. You can use dedicated ports or ports that also have data traffic.

Configuring intercluster LIFs on dedicated ports

You can configure intercluster LIFs on dedicated ports. Doing so typically increases the available bandwidth for replication traffic.

Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in "cluster01":

```
cluster01::> network port show
```

(Mbps)	Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed	Admin/Oper

cluster01-01								
		e0a	Cluster	Cluster	up	1500	auto/1000	
		e0b	Cluster	Cluster	up	1500	auto/1000	
		e0c	Default	Default	up	1500	auto/1000	
		e0d	Default	Default	up	1500	auto/1000	
		e0e	Default	Default	up	1500	auto/1000	
		e0f	Default	Default	up	1500	auto/1000	
cluster01-02								
		e0a	Cluster	Cluster	up	1500	auto/1000	
		e0b	Cluster	Cluster	up	1500	auto/1000	
		e0c	Default	Default	up	1500	auto/1000	
		e0d	Default	Default	up	1500	auto/1000	
		e0e	Default	Default	up	1500	auto/1000	
		e0f	Default	Default	up	1500	auto/1000	

2. Determine which ports are available to dedicate to intercluster communication:

```
network interface show -fields home-port,curr-port
```

For complete command syntax, see the man page.

The following example shows that ports "e0e" and "e0f" have not been assigned LIFs:

```

cluster01::> network interface show -fields home-port,curr-port
vserver lif                home-port curr-port
-----
Cluster cluster01-01_clus1 e0a       e0a
Cluster cluster01-01_clus2 e0b       e0b
Cluster cluster01-02_clus1 e0a       e0a
Cluster cluster01-02_clus2 e0b       e0b
cluster01
      cluster_mgmt         e0c       e0c
cluster01
      cluster01-01_mgmt1   e0c       e0c
cluster01
      cluster01-02_mgmt1   e0c       e0c

```

3. Create a failover group for the dedicated ports:

```

network interface failover-groups create -vserver <system_svm> -failover-group
<failover_group> -targets <physical_or_logical_ports>

```

The following example assigns ports "e0e" and "e0f" to failover group "intercluster01" on system "SVMcluster01":

```

cluster01::> network interface failover-groups create -vserver cluster01
-failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f

```

4. Verify that the failover group was created:

```

network interface failover-groups show

```

For complete command syntax, see the man page.

```

cluster01::> network interface failover-groups show

```

Vserver	Group	Failover Targets
Cluster	Cluster	cluster01-01:e0a, cluster01-01:e0b, cluster01-02:e0a, cluster01-02:e0b
cluster01	Default	cluster01-01:e0c, cluster01-01:e0d, cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f
	intercluster01	cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

In ONTAP 9.6 and later, run:

```

network interface create -vserver <system_svm> -lif <lif_name> -service
-policy default-intercluster -home-node <node_name> -home-port <port_name>
-address <port_ip_address> -netmask <netmask_address> -failover-group
<failover_group>

```

In ONTAP 9.5 and earlier, run:

```

network interface create -vserver <system_svm> -lif <lif_name> -role
intercluster -home-node <node_name> -home-port <port_name> -address
<port_ip_address> -netmask <netmask_address> -failover-group
<failover_group>

```

For complete command syntax, see the man page.

The following example creates intercluster LIFs "cluster01_icl01" and "cluster01_icl02" in failover group "intercluster01":

```

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0 -failover-group intercluster01

```

6. Verify that the intercluster LIFs were created:

In ONTAP 9.6 and later, run:

```
network interface show -service-policy default-intercluster
```

In ONTAP 9.5 and earlier, run:

```
network interface show -role intercluster
```

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-intercluster

```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
cluster01	cluster01_icl01	up/up	192.168.1.201/24	cluster01-01	e0e
true	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0f
true					

7. Verify that the intercluster LIFs are redundant:

In ONTAP 9.6 and later, run:

```
network interface show -service-policy default-intercluster -failover
```

In ONTAP 9.5 and earlier, run:

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01_icl01", and "cluster01_icl02" on the "SVMe0e" port will fail over to the "e0f" port.

```
cluster01::> network interface show -service-policy default-intercluster
-failover
          Logical          Home          Failover          Failover
Vserver  Interface          Node:Port          Policy            Group
-----
cluster01
          cluster01_icl01 cluster01-01:e0e   local-only
intercluster01
                                Failover Targets: cluster01-01:e0e,
                                cluster01-01:e0f
          cluster01_icl02 cluster01-02:e0e   local-only
intercluster01
                                Failover Targets: cluster01-02:e0e,
                                cluster01-02:e0f
```

Related information

[Considerations when using dedicated ports](#)

Configuring intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in "cluster01":

```
cluster01::> network port show
```

(Mbps)							Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	

cluster01-01							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	
cluster01-02							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	

2. Create intercluster LIFs on the system SVM:

In ONTAP 9.6 and later, run:

```
network interface create -vserver <system_svm> -lif <lif_name> -service
-policy default-intercluster -home-node <node_name> -home-port <port_name>
-address <port_ip_address> -netmask <netmask>
```

In ONTAP 9.5 and earlier, run:

```
network interface create -vserver <system_svm> -lif <lif_name> -role
intercluster -home-node <node_name> -home-port <port_name> -address
<port_ip_address> -netmask <netmask>
```

For complete command syntax, see the man page.

The following example creates intercluster LIFs "cluster01_icl01" and "cluster01_icl02":

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0c
-address 192.168.1.201
-netmask 255.255.255.0

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0c
-address 192.168.1.202
-netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

In ONTAP 9.6 and later, run:

```
network interface show -service-policy default-intercluster
```

In ONTAP 9.5 and earlier, run:

```
network interface show -role intercluster
```

For complete command syntax, see the man page.

```
cluster01::> network interface show -service-policy default-intercluster
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
cluster01	cluster01_icl01	up/up	192.168.1.201/24	cluster01-01	e0c
true	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0c
true					

4. Verify that the intercluster LIFs are redundant:

In ONTAP 9.6 and later, run:

```
network interface show -service-policy default-intercluster -failover
```

In ONTAP 9.5 and earlier, run:

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that intercluster LIFs "cluster01_icl01" and "cluster01_icl02" on the "e0c" port will fail over to the "e0d" port.

```

cluster01::> network interface show -service-policy default-intercluster
-failover

```

Vserver	Logical Interface	Home Node:Port	Failover Policy	Failover Group
cluster01	cluster01_icl01	cluster01-01:e0c	local-only	
	192.168.1.201/24			
			Failover Targets: cluster01-01:e0c, cluster01-01:e0d	
	cluster01_icl02	cluster01-02:e0c	local-only	
	192.168.1.201/24			
			Failover Targets: cluster01-02:e0c, cluster01-02:e0d	

Related information

[Considerations when sharing data ports](#)

Creating a cluster peer relationship

You can use the `cluster peer create` command to create a peer relationship between a local and remote cluster. After the peer relationship has been created, you can run `cluster peer create` on the remote cluster to authenticate it to the local cluster.

About this task

- You must have created intercluster LIFs on every node in the clusters that are being peered.
- The clusters must be running ONTAP 9.3 or later.

Steps

1. On the destination cluster, create a peer relationship with the source cluster:

```

cluster peer create -generate-passphrase -offer-expiration <MM/DD/YYYY
HH:MM:SS|1...7days|1...168hours> -peer-addr <peer_lif_ip_addresses> -ipspace
<ipspace>

```

If you specify both `-generate-passphrase` and `-peer-addr`, only the cluster whose intercluster LIFs are specified in `-peer-addr` can use the generated password.

You can ignore the `-ipspace` option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship on an unspecified remote cluster:

```
cluster02::> cluster peer create -generate-passphrase -offer-expiration
2days
```

```
                Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
                Expiration Time: 6/7/2017 08:16:10 EST
Initial Allowed Vserver Peers: -
                Intercluster LIF IP: 192.140.112.101
                Peer Cluster Name: Clus_7ShR (temporary generated)
```

Warning: make a note of the passphrase - it cannot be displayed again.

2. On the source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addr <peer_lif_ip_addresses> -ip-space <ip-space>
```

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses "192.140.112.101" and "192.140.112.102":

```
cluster01::> cluster peer create -peer-addr
192.140.112.101,192.140.112.102
```

Notice: Use a generated passphrase or choose a passphrase of 8 or more characters.

To ensure the authenticity of the peering relationship, use a phrase or sequence of characters that would be hard to guess.

```
Enter the passphrase:
Confirm the passphrase:
```

```
Clusters cluster02 and cluster01 are peered.
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

```
cluster01::> cluster peer show -instance

Peer Cluster Name: cluster02
Cluster UUID: b07036f2-7d1c-11f0-bedb-
d039ea48b059
Remote Intercluster Addresses: 192.140.112.101,
192.140.112.102
Availability of the Remote Cluster: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.140.112.101,
192.140.112.102
Cluster Serial Number: 1-80-123456
Remote Cluster Nodes: cluster02-01, cluster02-02,
Remote Cluster Health: true
Unreachable Local Nodes: -
Operation Timeout (seconds): 60
Address Family of Relationship: ipv4
Authentication Status Administrative: use-authentication
Authentication Status Operational: ok
Timeout for RPC Connect: 10
Timeout for Update Pings: 5
Last Update Time: 10/9/2025 10:15:29
IPspace for the Relationship: Default
Proposed Setting for Encryption of Inter-Cluster Communication: -
Encryption Protocol For Inter-Cluster Communication: tls-psk
Algorithm By Which the PSK Was Derived: jpake
```

4. Check the connectivity and status of the nodes in the peer relationship:

```
cluster peer health show
```

```

cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
-----
cluster01-01
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
cluster01-02
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true

```

Creating the DR group

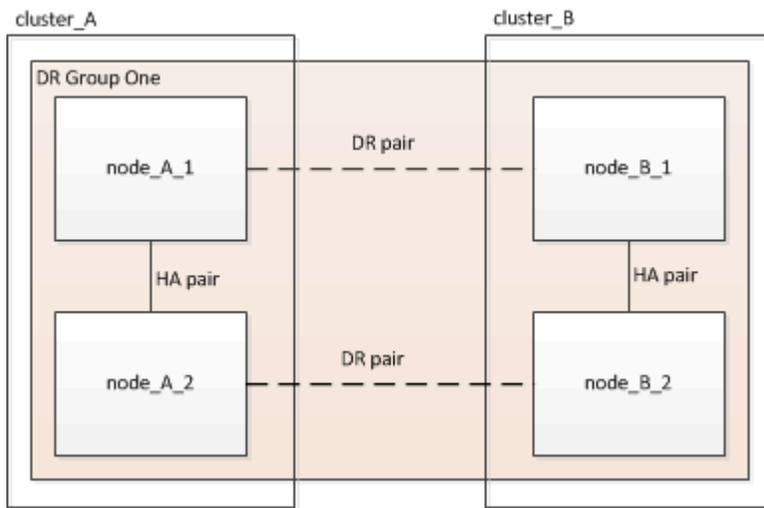
You must create the disaster recovery (DR) group relationships between the clusters.

About this task

You perform this procedure on one of the clusters in the MetroCluster configuration to create the DR relationships between the nodes in both clusters.



The DR relationships cannot be changed after the DR groups are created.



Steps

1. Verify that the nodes are ready for creation of the DR group by entering the following command on each

node:

```
metrocluster configuration-settings show-status
```

The command output should show that the nodes are ready:

```
cluster_A::> metrocluster configuration-settings show-status
Cluster                Node                Configuration Settings Status
-----
cluster_A              node_A_1            ready for DR group create
                       node_A_2            ready for DR group create
2 entries were displayed.
```

```
cluster_B::> metrocluster configuration-settings show-status
Cluster                Node                Configuration Settings Status
-----
cluster_B              node_B_1            ready for DR group create
                       node_B_2            ready for DR group create
2 entries were displayed.
```

2. Create the DR group:

```
metrocluster configuration-settings dr-group create -partner-cluster
<partner_cluster_name> -local-node <local_node_name> -remote-node
<remote_node_name>
```

This command is issued only once. It does not need to be repeated on the partner cluster. In the command, you specify the name of the remote cluster and the name of one local node and one node on the partner cluster.

The two nodes you specify are configured as DR partners and the other two nodes (which are not specified in the command) are configured as the second DR pair in the DR group. These relationships cannot be changed after you enter this command.

The following command creates these DR pairs:

- node_A_1 and node_B_1
- node_A_2 and node_B_2

```
Cluster_A::> metrocluster configuration-settings dr-group create
-partner-cluster cluster_B -local-node node_A_1 -remote-node node_B_1
[Job 27] Job succeeded: DR Group Create is successful.
```

Configuring and connecting the MetroCluster IP interfaces

You must configure the MetroCluster IP interfaces that are used for replication of each node's storage and nonvolatile cache. You then establish the connections using the MetroCluster IP interfaces. This creates iSCSI connections for storage replication.



The MetroCluster IP and connected switch ports do not come online until after you create the MetroCluster IP interfaces.

About this task

- You must create two interfaces for each node. The interfaces must be associated with the VLANs defined in the MetroCluster RCF file.
- You must create all MetroCluster IP interface "A" ports in the same VLAN and all MetroCluster IP interface "B" ports in the other VLAN. Refer to [Considerations for MetroCluster IP configuration](#).
- Beginning with ONTAP 9.9.1, if you are using a layer 3 configuration, you must also specify the `-gateway` parameter when creating MetroCluster IP interfaces. Refer to [Considerations for layer 3 wide-area networks](#).

Certain platforms use a VLAN for the MetroCluster IP interface. By default, each of the two ports use a different VLAN: 10 and 20.

If supported, you can also specify a different (non-default) VLAN higher than 100 (between 101 and 4095) using the `-vlan-id` parameter in the `metrocluster configuration-settings interface create` command.

The following platforms do **not** support the `-vlan-id` parameter:

- FAS8200 and AFF A300
- AFF A320
- FAS9000 and AFF A700
- AFF C800, ASA C800, AFF A800 and ASA A800

All other platforms support the `-vlan-id` parameter.

The default and valid VLAN assignments depend on whether the platform supports the `-vlan-id` parameter:

Platforms that support `-vlan-id`

Default VLAN:

- When the `-vlan-id` parameter is not specified, the interfaces are created with VLAN 10 for the "A" ports and VLAN 20 for the "B" ports.
- The VLAN specified must match the VLAN selected in the RCF.

Valid VLAN ranges:

- Default VLAN 10 and 20
- VLANs 101 and higher (between 101 and 4095)

Platforms that do not support `-vlan-id`

Default VLAN:

- Not applicable. The interface does not require a VLAN to be specified on the MetroCluster interface. The switch port defines the VLAN that is used.

Valid VLAN ranges:

- All VLANs not explicitly excluded when generating the RCF. The RCF alerts you if the VLAN is invalid.

- The physical ports used by the MetroCluster IP interfaces depends on the platform model. Refer to [Cable the MetroCluster IP switches](#) for the port usage for your system.
- The following IP addresses and subnets are used in the examples:

Node	Interface	IP address	Subnet
node_A_1	MetroCluster IP interface 1	10.1.1.1	10.1.1/24
	MetroCluster IP interface 2	10.1.2.1	10.1.2/24
node_A_2	MetroCluster IP interface 1	10.1.1.2	10.1.1/24
	MetroCluster IP interface 2	10.1.2.2	10.1.2/24
node_B_1	MetroCluster IP interface 1	10.1.1.3	10.1.1/24
	MetroCluster IP interface 2	10.1.2.3	10.1.2/24

node_B_2	MetroCluster IP interface 1	10.1.1.4	10.1.1/24
	MetroCluster IP interface 2	10.1.2.4	10.1.2/24

- This procedure uses the following examples:

The ports for an AFF A700 or a FAS9000 system (e5a and e5b).

The ports for an AFF A220 system to show how to use the `-vlan-id` parameter on a supported platform.

Configure the interfaces on the correct ports for your platform model.

Steps

1. Confirm that each node has disk automatic assignment enabled:

```
storage disk option show
```

Disk automatic assignment will assign pool 0 and pool 1 disks on a shelf-by-shelf basis.

The Auto Assign column indicates whether disk automatic assignment is enabled.

```

Node           BKg. FW. Upd.  Auto Copy  Auto Assign  Auto Assign Policy
-----
node_A_1              on          on           on           default
node_A_2              on          on           on           default
2 entries were displayed.

```

2. Verify you can create MetroCluster IP interfaces on the nodes:

```
metrocluster configuration-settings show-status
```

All nodes should be ready:

```

Cluster       Node           Configuration Settings Status
-----
cluster_A
node_A_1      ready for interface create
node_A_2      ready for interface create
cluster_B
node_B_1      ready for interface create
node_B_2      ready for interface create
4 entries were displayed.

```

3. Create the interfaces on node_A_1.

- a. Configure the interface on port "e5a" on "node_A_1":



Do not use 169.254.17.x or 169.254.18.x IP addresses when you create MetroCluster IP interfaces to avoid conflicts with system auto-generated interface IP addresses in the same range.

```
metrocluster configuration-settings interface create -cluster-name
<cluster_name> -home-node <node_name> -home-port e5a -address <ip_address>
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5a" on "node_A_1" with IP address "10.1.1.1":

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_1 -home-port e5a -address
10.1.1.1 -netmask 255.255.255.0
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```

On platform models that support VLANs for the MetroCluster IP interface, you can include the `-vlan-id` parameter if you don't want to use the default VLAN IDs. The following example shows the command for an AFF A220 system with a VLAN ID of 120:

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_2 -home-port e0a -address
10.1.1.2 -netmask 255.255.255.0 -vlan-id 120
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```

- b. Configure the interface on port "e5b" on "node_A_1":

```
metrocluster configuration-settings interface create -cluster-name
<cluster_name> -home-node <node_name> -home-port e5b -address <ip_address>
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5b" on "node_A_1" with IP address "10.1.2.1":

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_1 -home-port e5b -address
10.1.2.1 -netmask 255.255.255.0
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```



You can verify that these interfaces are present using the `metrocluster configuration-settings interface show` command.

4. Create the interfaces on node_A_2.

a. Configure the interface on port "e5a" on "node_A_2":

```
metrocluster configuration-settings interface create -cluster-name
<cluster_name> -home-node <node_name> -home-port e5a -address <ip_address>
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5a" on "node_A_2" with IP address "10.1.1.2":

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_2 -home-port e5a -address
10.1.1.2 -netmask 255.255.255.0
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```

b. Configure the interface on port "e5b" on "node_A_2":

```
metrocluster configuration-settings interface create -cluster-name
<cluster_name> -home-node <node_name> -home-port e5b -address <ip_address>
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5b" on "node_A_2" with IP address "10.1.2.2":

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_2 -home-port e5b -address
10.1.2.2 -netmask 255.255.255.0
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```

On platform models that support VLANs for the MetroCluster IP interface, you can include the `-vlan -id` parameter if you don't want to use the default VLAN IDs. The following example shows the command for an AFF A220 system with a VLAN ID of 220:

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_2 -home-port e0b -address
10.1.2.2 -netmask 255.255.255.0 -vlan-id 220
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```

5. Create the interfaces on "node_B_1".

- a. Configure the interface on port "e5a" on "node_B_1":

```
metrocluster configuration-settings interface create -cluster-name  
<cluster_name> -home-node <node_name> -home-port e5a -address <ip_address>  
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5a" on "node_B_1" with IP address "10.1.1.3":

```
cluster_A::> metrocluster configuration-settings interface create  
-cluster-name cluster_B -home-node node_B_1 -home-port e5a -address  
10.1.1.3 -netmask 255.255.255.0  
[Job 28] Job succeeded: Interface Create is successful.cluster_B::>
```

- b. Configure the interface on port "e5b" on "node_B_1":

```
metrocluster configuration-settings interface create -cluster-name  
<cluster_name> -home-node <node_name> -home-port e5b -address <ip_address>  
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5b" on "node_B_1" with IP address "10.1.2.3":

```
cluster_A::> metrocluster configuration-settings interface create  
-cluster-name cluster_B -home-node node_B_1 -home-port e5b -address  
10.1.2.3 -netmask 255.255.255.0  
[Job 28] Job succeeded: Interface Create is successful.cluster_B::>
```

6. Create the interfaces on "node_B_2".

- a. Configure the interface on port e5a on node_B_2:

```
metrocluster configuration-settings interface create -cluster-name  
<cluster_name> -home-node <node_name> -home-port e5a -address <ip_address>  
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5a" on "node_B_2" with IP address "10.1.1.4":

```
cluster_B::> metrocluster configuration-settings interface create  
-cluster-name cluster_B -home-node node_B_2 -home-port e5a -address  
10.1.1.4 -netmask 255.255.255.0  
[Job 28] Job succeeded: Interface Create is successful.cluster_A::>
```

- b. Configure the interface on port "e5b" on "node_B_2":

```
metrocluster configuration-settings interface create -cluster-name  
<cluster_name> -home-node <node_name> -home-port e5b -address <ip_address>
```

```
-netmask <netmask>
```

The following example shows the creation of the interface on port "e5b" on "node_B_2" with IP address "10.1.2.4":

```
cluster_B::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_2 -home-port e5b -address
10.1.2.4 -netmask 255.255.255.0
[Job 28] Job succeeded: Interface Create is successful.
cluster_A::>
```

7. Verify that the interfaces have been configured:

```
metrocluster configuration-settings interface show
```

The following example shows that the configuration state for each interface is completed.

```
cluster_A::> metrocluster configuration-settings interface show
DR
Group Cluster Node      Network Address Netmask      Gateway  State
-----
1      cluster_A  node_A_1
      Home Port: e5a
      10.1.1.1      255.255.255.0  -        completed
      Home Port: e5b
      10.1.2.1      255.255.255.0  -        completed
      node_A_2
      Home Port: e5a
      10.1.1.2      255.255.255.0  -        completed
      Home Port: e5b
      10.1.2.2      255.255.255.0  -        completed
      cluster_B node_B_1
      Home Port: e5a
      10.1.1.3      255.255.255.0  -        completed
      Home Port: e5b
      10.1.2.3      255.255.255.0  -        completed
      node_B_2
      Home Port: e5a
      10.1.1.4      255.255.255.0  -        completed
      Home Port: e5b
      10.1.2.4      255.255.255.0  -        completed
8 entries were displayed.
cluster_A::>
```

8. Verify that the nodes are ready to connect the MetroCluster interfaces:

```
metrocluster configuration-settings show-status
```

The following example shows all nodes in the "ready for connection" state:

```
Cluster      Node      Configuration Settings Status
-----
cluster_A
            node_A_1  ready for connection connect
            node_A_2  ready for connection connect
cluster_B
            node_B_1  ready for connection connect
            node_B_2  ready for connection connect
4 entries were displayed.
```

9. Establish the connections:

```
metrocluster configuration-settings connection connect
```

If you are running a version earlier than ONTAP 9.10.1, the IP addresses cannot be changed after you issue this command.

The following example shows cluster_A is successfully connected:

```
cluster_A::> metrocluster configuration-settings connection connect
[Job 53] Job succeeded: Connect is successful.
cluster_A::>
```

10. Verify that the connections have been established:

```
metrocluster configuration-settings show-status
```

The configuration settings status for all nodes should be completed:

```
Cluster      Node      Configuration Settings Status
-----
cluster_A
            node_A_1  completed
            node_A_2  completed
cluster_B
            node_B_1  completed
            node_B_2  completed
4 entries were displayed.
```

11. Verify that the iSCSI connections have been established:

a. Change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when you are prompted to continue into advanced mode and you see the advanced mode prompt (`*>`).

b. Display the connections:

```
storage iscsi-initiator show
```

On systems running ONTAP 9.5, there are eight MetroCluster IP initiators on each cluster that should appear in the output.

On systems running ONTAP 9.4 and earlier, there are four MetroCluster IP initiators on each cluster that should appear in the output.

The following example shows the eight MetroCluster IP initiators on a cluster running ONTAP 9.5:

```
cluster_A::*> storage iscsi-initiator show
Node Type Label      Target Portal      Target Name
Admin/Op
-----
-----

cluster_A-01
  dr_auxiliary
    mccip-aux-a-initiator
      10.227.16.113:65200      prod506.com.company:abab44
up/up
    mccip-aux-a-initiator2
      10.227.16.113:65200      prod507.com.company:abab44
up/up
    mccip-aux-b-initiator
      10.227.95.166:65200      prod506.com.company:abab44
up/up
    mccip-aux-b-initiator2
      10.227.95.166:65200      prod507.com.company:abab44
up/up
  dr_partner
    mccip-pri-a-initiator
      10.227.16.112:65200      prod506.com.company:cdcd88
up/up
    mccip-pri-a-initiator2
      10.227.16.112:65200      prod507.com.company:cdcd88
up/up
    mccip-pri-b-initiator
      10.227.95.165:65200      prod506.com.company:cdcd88
```

```

up/up
    mccip-pri-b-initiator2
        10.227.95.165:65200      prod507.com.company:cdcd88
up/up
cluster_A-02
  dr_auxiliary
    mccip-aux-a-initiator
        10.227.16.112:65200    prod506.com.company:cdcd88
up/up
    mccip-aux-a-initiator2
        10.227.16.112:65200    prod507.com.company:cdcd88
up/up
    mccip-aux-b-initiator
        10.227.95.165:65200    prod506.com.company:cdcd88
up/up
    mccip-aux-b-initiator2
        10.227.95.165:65200    prod507.com.company:cdcd88
up/up
  dr_partner
    mccip-pri-a-initiator
        10.227.16.113:65200    prod506.com.company:abab44
up/up
    mccip-pri-a-initiator2
        10.227.16.113:65200    prod507.com.company:abab44
up/up
    mccip-pri-b-initiator
        10.227.95.166:65200    prod506.com.company:abab44
up/up
    mccip-pri-b-initiator2
        10.227.95.166:65200    prod507.com.company:abab44
up/up
16 entries were displayed.

```

c. Return to the admin privilege level:

```
set -privilege admin
```

12. Verify that the nodes are ready for final implementation of the MetroCluster configuration:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR
Group Cluster Node          Configuration  DR
State          Mirroring Mode
-----
-      cluster_A
           node_A_1          ready to configure -    -
           node_A_2          ready to configure -    -
2 entries were displayed.
cluster_A::>

```

```

cluster_B::> metrocluster node show
DR
Group Cluster Node          Configuration  DR
State          Mirroring Mode
-----
-      cluster_B
           node_B_1          ready to configure -    -
           node_B_2          ready to configure -    -
2 entries were displayed.
cluster_B::>

```

Verifying or manually performing pool 1 drives assignment

Depending on the storage configuration, you must either verify pool 1 drive assignment or manually assign drives to pool 1 for each node in the MetroCluster IP configuration. The procedure you use depends on the version of ONTAP you are using.

Configuration type	Procedure
The systems meet the requirements for automatic drive assignment or, if running ONTAP 9.3, were received from the factory.	Verifying disk assignment for pool 1 disks
The configuration includes either three shelves, or, if it contains more than four shelves, has an uneven multiple of four shelves (for example, seven shelves), and is running ONTAP 9.5.	Manually assigning drives for pool 1 (ONTAP 9.4 or later)
The configuration does not include four storage shelves per site and is running ONTAP 9.4	Manually assigning drives for pool 1 (ONTAP 9.4 or later)
The systems were not received from the factory and are running ONTAP 9.3 Systems received from the factory are pre-configured with assigned drives.	Manually assigning disks for pool 1 (ONTAP 9.3)

Verifying disk assignment for pool 1 disks

You must verify that the remote disks are visible to the nodes and have been assigned correctly.

Before you begin

You must wait at least ten minutes for disk auto-assignment to complete after the MetroCluster IP interfaces and connections were created with the `metrocluster configuration-settings connection connect` command.

Command output will show disk names in the form: `node-name:0m.i1.0L1`

Considerations for automatic drive assignment and ADP systems in ONTAP 9.4 and later

Steps

1. Verify pool 1 disks are auto-assigned:

```
disk show
```

The following output shows the output for an AFF A800 system with no external shelves.

Drive autoassignment has assigned one quarter (8 drives) to "node_A_1" and one quarter to "node_A_2". The remaining drives will be remote (pool 1) disks for "node_B_1" and "node_B_2".

```
cluster_B::> disk show -host-adapter 0m -owner node_B_2
          Usable      Disk          Container  Container
Disk      Size      Shelf Bay Type      Type      Name
Owner
-----
node_B_2:0m.i0.2L4  894.0GB    0     29  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.2L10 894.0GB    0     25  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.3L3   894.0GB    0     28  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.3L9   894.0GB    0     24  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.3L11 894.0GB    0     26  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.3L12 894.0GB    0     27  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.3L15 894.0GB    0     30  SSD-NVM  shared    -
node_B_2
node_B_2:0m.i0.3L16 894.0GB    0     31  SSD-NVM  shared    -
node_B_2
8 entries were displayed.

cluster_B::> disk show -host-adapter 0m -owner node_B_1
          Usable      Disk          Container  Container
```

```

Disk          Size      Shelf Bay Type      Type      Name
Owner
-----
-----
node_B_1:0m.i2.3L19 1.75TB    0      42  SSD-NVM  shared    -
node_B_1
node_B_1:0m.i2.3L20 1.75TB    0      43  SSD-NVM  spare     Pool1
node_B_1
node_B_1:0m.i2.3L23 1.75TB    0      40  SSD-NVM  shared    -
node_B_1
node_B_1:0m.i2.3L24 1.75TB    0      41  SSD-NVM  spare     Pool1
node_B_1
node_B_1:0m.i2.3L29 1.75TB    0      36  SSD-NVM  shared    -
node_B_1
node_B_1:0m.i2.3L30 1.75TB    0      37  SSD-NVM  shared    -
node_B_1
node_B_1:0m.i2.3L31 1.75TB    0      38  SSD-NVM  shared    -
node_B_1
node_B_1:0m.i2.3L32 1.75TB    0      39  SSD-NVM  shared    -
node_B_1
8 entries were displayed.

```

```
cluster_B::> disk show
```

```

          Usable      Disk          Container      Container
Disk      Size      Shelf Bay Type      Type      Name
Owner
-----
-----
node_B_1:0m.i1.0L6  1.75TB    0      1   SSD-NVM  shared    -
node_A_2
node_B_1:0m.i1.0L8  1.75TB    0      3   SSD-NVM  shared    -
node_A_2
node_B_1:0m.i1.0L17 1.75TB    0     18   SSD-NVM  shared    -
node_A_1
node_B_1:0m.i1.0L22 1.75TB    0     17   SSD-NVM  shared - node_A_1
node_B_1:0m.i1.0L25 1.75TB    0     12   SSD-NVM  shared - node_A_1
node_B_1:0m.i1.2L2  1.75TB    0      5   SSD-NVM  shared - node_A_2
node_B_1:0m.i1.2L7  1.75TB    0      2   SSD-NVM  shared - node_A_2
node_B_1:0m.i1.2L14 1.75TB    0      7   SSD-NVM  shared - node_A_2
node_B_1:0m.i1.2L21 1.75TB    0     16   SSD-NVM  shared - node_A_1
node_B_1:0m.i1.2L27 1.75TB    0     14   SSD-NVM  shared - node_A_1
node_B_1:0m.i1.2L28 1.75TB    0     15   SSD-NVM  shared - node_A_1
node_B_1:0m.i2.1L1  1.75TB    0      4   SSD-NVM  shared - node_A_2
node_B_1:0m.i2.1L5  1.75TB    0      0   SSD-NVM  shared - node_A_2
node_B_1:0m.i2.1L13 1.75TB    0      6   SSD-NVM  shared - node_A_2
node_B_1:0m.i2.1L18 1.75TB    0     19   SSD-NVM  shared - node_A_1

```

```

node_B_1:0m.i2.1L26 1.75TB 0 13 SSD-NVM shared - node_A_1
node_B_1:0m.i2.3L19 1.75TB 0 42 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L20 1.75TB 0 43 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L23 1.75TB 0 40 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L24 1.75TB 0 41 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L29 1.75TB 0 36 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L30 1.75TB 0 37 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L31 1.75TB 0 38 SSD-NVM shared - node_B_1
node_B_1:0m.i2.3L32 1.75TB 0 39 SSD-NVM shared - node_B_1
node_B_1:0n.12 1.75TB 0 12 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.13 1.75TB 0 13 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.14 1.75TB 0 14 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.15 1.75TB 0 15 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.16 1.75TB 0 16 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.17 1.75TB 0 17 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.18 1.75TB 0 18 SSD-NVM shared aggr0 node_B_1
node_B_1:0n.19 1.75TB 0 19 SSD-NVM shared - node_B_1
node_B_1:0n.24 894.0GB 0 24 SSD-NVM shared - node_A_2
node_B_1:0n.25 894.0GB 0 25 SSD-NVM shared - node_A_2
node_B_1:0n.26 894.0GB 0 26 SSD-NVM shared - node_A_2
node_B_1:0n.27 894.0GB 0 27 SSD-NVM shared - node_A_2
node_B_1:0n.28 894.0GB 0 28 SSD-NVM shared - node_A_2
node_B_1:0n.29 894.0GB 0 29 SSD-NVM shared - node_A_2
node_B_1:0n.30 894.0GB 0 30 SSD-NVM shared - node_A_2
node_B_1:0n.31 894.0GB 0 31 SSD-NVM shared - node_A_2
node_B_1:0n.36 1.75TB 0 36 SSD-NVM shared - node_A_1
node_B_1:0n.37 1.75TB 0 37 SSD-NVM shared - node_A_1
node_B_1:0n.38 1.75TB 0 38 SSD-NVM shared - node_A_1
node_B_1:0n.39 1.75TB 0 39 SSD-NVM shared - node_A_1
node_B_1:0n.40 1.75TB 0 40 SSD-NVM shared - node_A_1
node_B_1:0n.41 1.75TB 0 41 SSD-NVM shared - node_A_1
node_B_1:0n.42 1.75TB 0 42 SSD-NVM shared - node_A_1
node_B_1:0n.43 1.75TB 0 43 SSD-NVM shared - node_A_1
node_B_2:0m.i0.2L4 894.0GB 0 29 SSD-NVM shared - node_B_2
node_B_2:0m.i0.2L10 894.0GB 0 25 SSD-NVM shared - node_B_2
node_B_2:0m.i0.3L3 894.0GB 0 28 SSD-NVM shared - node_B_2
node_B_2:0m.i0.3L9 894.0GB 0 24 SSD-NVM shared - node_B_2
node_B_2:0m.i0.3L11 894.0GB 0 26 SSD-NVM shared - node_B_2
node_B_2:0m.i0.3L12 894.0GB 0 27 SSD-NVM shared - node_B_2
node_B_2:0m.i0.3L15 894.0GB 0 30 SSD-NVM shared - node_B_2
node_B_2:0m.i0.3L16 894.0GB 0 31 SSD-NVM shared - node_B_2
node_B_2:0n.0 1.75TB 0 0 SSD-NVM shared aggr0_rha12_b1_cm_02_0
node_B_2
node_B_2:0n.1 1.75TB 0 1 SSD-NVM shared aggr0_rha12_b1_cm_02_0 node_B_2
node_B_2:0n.2 1.75TB 0 2 SSD-NVM shared aggr0_rha12_b1_cm_02_0 node_B_2
node_B_2:0n.3 1.75TB 0 3 SSD-NVM shared aggr0_rha12_b1_cm_02_0 node_B_2

```

```

node_B_2:0n.4 1.75TB 0 4 SSD-NVM shared aggr0_rha12_b1_cm_02_0 node_B_2
node_B_2:0n.5 1.75TB 0 5 SSD-NVM shared aggr0_rha12_b1_cm_02_0 node_B_2
node_B_2:0n.6 1.75TB 0 6 SSD-NVM shared aggr0_rha12_b1_cm_02_0 node_B_2
node_B_2:0n.7 1.75TB 0 7 SSD-NVM shared - node_B_2
64 entries were displayed.

```

```
cluster_B::>
```

```
cluster_A::> disk show
```

```
Usable Disk Container Container
```

```
Disk Size Shelf Bay Type Type Name Owner
```

```

-----
-----
node_A_1:0m.i1.0L2 1.75TB 0 5 SSD-NVM shared - node_B_2
node_A_1:0m.i1.0L8 1.75TB 0 3 SSD-NVM shared - node_B_2
node_A_1:0m.i1.0L18 1.75TB 0 19 SSD-NVM shared - node_B_1
node_A_1:0m.i1.0L25 1.75TB 0 12 SSD-NVM shared - node_B_1
node_A_1:0m.i1.0L27 1.75TB 0 14 SSD-NVM shared - node_B_1
node_A_1:0m.i1.2L1 1.75TB 0 4 SSD-NVM shared - node_B_2
node_A_1:0m.i1.2L6 1.75TB 0 1 SSD-NVM shared - node_B_2
node_A_1:0m.i1.2L7 1.75TB 0 2 SSD-NVM shared - node_B_2
node_A_1:0m.i1.2L14 1.75TB 0 7 SSD-NVM shared - node_B_2
node_A_1:0m.i1.2L17 1.75TB 0 18 SSD-NVM shared - node_B_1
node_A_1:0m.i1.2L22 1.75TB 0 17 SSD-NVM shared - node_B_1
node_A_1:0m.i2.1L5 1.75TB 0 0 SSD-NVM shared - node_B_2
node_A_1:0m.i2.1L13 1.75TB 0 6 SSD-NVM shared - node_B_2
node_A_1:0m.i2.1L21 1.75TB 0 16 SSD-NVM shared - node_B_1
node_A_1:0m.i2.1L26 1.75TB 0 13 SSD-NVM shared - node_B_1
node_A_1:0m.i2.1L28 1.75TB 0 15 SSD-NVM shared - node_B_1
node_A_1:0m.i2.3L19 1.75TB 0 42 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L20 1.75TB 0 43 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L23 1.75TB 0 40 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L24 1.75TB 0 41 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L29 1.75TB 0 36 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L30 1.75TB 0 37 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L31 1.75TB 0 38 SSD-NVM shared - node_A_1
node_A_1:0m.i2.3L32 1.75TB 0 39 SSD-NVM shared - node_A_1
node_A_1:0n.12 1.75TB 0 12 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.13 1.75TB 0 13 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.14 1.75TB 0 14 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.15 1.75TB 0 15 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.16 1.75TB 0 16 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.17 1.75TB 0 17 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.18 1.75TB 0 18 SSD-NVM shared aggr0 node_A_1
node_A_1:0n.19 1.75TB 0 19 SSD-NVM shared - node_A_1

```

```

node_A_1:0n.24 894.0GB 0 24 SSD-NVM shared - node_B_2
node_A_1:0n.25 894.0GB 0 25 SSD-NVM shared - node_B_2
node_A_1:0n.26 894.0GB 0 26 SSD-NVM shared - node_B_2
node_A_1:0n.27 894.0GB 0 27 SSD-NVM shared - node_B_2
node_A_1:0n.28 894.0GB 0 28 SSD-NVM shared - node_B_2
node_A_1:0n.29 894.0GB 0 29 SSD-NVM shared - node_B_2
node_A_1:0n.30 894.0GB 0 30 SSD-NVM shared - node_B_2
node_A_1:0n.31 894.0GB 0 31 SSD-NVM shared - node_B_2
node_A_1:0n.36 1.75TB 0 36 SSD-NVM shared - node_B_1
node_A_1:0n.37 1.75TB 0 37 SSD-NVM shared - node_B_1
node_A_1:0n.38 1.75TB 0 38 SSD-NVM shared - node_B_1
node_A_1:0n.39 1.75TB 0 39 SSD-NVM shared - node_B_1
node_A_1:0n.40 1.75TB 0 40 SSD-NVM shared - node_B_1
node_A_1:0n.41 1.75TB 0 41 SSD-NVM shared - node_B_1
node_A_1:0n.42 1.75TB 0 42 SSD-NVM shared - node_B_1
node_A_1:0n.43 1.75TB 0 43 SSD-NVM shared - node_B_1
node_A_2:0m.i2.3L3 894.0GB 0 28 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L4 894.0GB 0 29 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L9 894.0GB 0 24 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L10 894.0GB 0 25 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L11 894.0GB 0 26 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L12 894.0GB 0 27 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L15 894.0GB 0 30 SSD-NVM shared - node_A_2
node_A_2:0m.i2.3L16 894.0GB 0 31 SSD-NVM shared - node_A_2
node_A_2:0n.0 1.75TB 0 0 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.1 1.75TB 0 1 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.2 1.75TB 0 2 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.3 1.75TB 0 3 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.4 1.75TB 0 4 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.5 1.75TB 0 5 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.6 1.75TB 0 6 SSD-NVM shared aggr0_node_A_2_0 node_A_2
node_A_2:0n.7 1.75TB 0 7 SSD-NVM shared - node_A_2
64 entries were displayed.

```

```
cluster_A::>
```

Manually assigning drives for pool 1 (ONTAP 9.4 or later)

If the system was not preconfigured at the factory and does not meet the requirements for automatic drive assignment, you must manually assign the remote pool 1 drives.

About this task

This procedure applies to configurations running ONTAP 9.4 or later.

Details for determining whether your system requires manual disk assignment are included in [Considerations for automatic drive assignment and ADP systems in ONTAP 9.4 and later](#).

When the configuration includes only two external shelves per site, pool 1 drives for each site should be shared from the same shelf as shown in the following examples:

- node_A_1 is assigned drives in bays 0-11 on site_B-shelf_2 (remote)
- node_A_2 is assigned drives in bays 12-23 on site_B-shelf_2 (remote)

Steps

1. From each node in the MetroCluster IP configuration, assign remote drives to pool 1.

a. Display the list of unassigned drives:

```
disk show -host-adapter 0m -container-type unassigned
```

```
cluster_A::> disk show -host-adapter 0m -container-type unassigned
              Usable          Disk      Container  Container
Disk          Size Shelf Bay Type      Type      Name
Owner
-----
-----
6.23.0         -    23   0 SSD    unassigned -    -
6.23.1         -    23   1 SSD    unassigned -    -
.
.
.
node_A_2:0m.i1.2L51 -    21  14 SSD    unassigned -    -
node_A_2:0m.i1.2L64 -    21  10 SSD    unassigned -    -
.
.
.
48 entries were displayed.

cluster_A::>
```

b. Assign ownership of remote drives (0m) to pool 1 of the first node (for example, node_A_1):

```
disk assign -disk <disk-id> -pool 1 -owner <owner_node_name>
```

disk-id must identify a drive on a remote shelf of owner_node_name.

c. Confirm that the drives were assigned to pool 1:

```
disk show -host-adapter 0m -container-type unassigned
```



The iSCSI connection used to access the remote drives appears as device 0m.

The following output shows that the drives on shelf 23 were assigned because they no longer appear in the list of unassigned drives:

```

cluster_A::> disk show -host-adapter 0m -container-type unassigned
              Usable           Disk   Container  Container
Disk          Size Shelf Bay Type   Type       Name
Owner
-----
-----
node_A_2:0m.i1.2L51      -    21  14 SSD    unassigned -    -
node_A_2:0m.i1.2L64      -    21  10 SSD    unassigned -    -
.
.
.
node_A_2:0m.i2.1L90      -    21  19 SSD    unassigned -    -
24 entries were displayed.

cluster_A::>

```

- d. Repeat these steps to assign pool 1 drives to the second node on site A (for example, "node_A_2").
- e. Repeat these steps on site B.

Manually assigning disks for pool 1 (ONTAP 9.3)

If you have at least two disk shelves for each node, you use ONTAP's auto-assignment functionality to automatically assign the remote (pool1) disks.

Before you begin

You must first assign a disk on the shelf to pool 1. ONTAP then automatically assigns the rest of the disks on the shelf to the same pool.

About this task

This procedure applies to configurations running ONTAP 9.3.

This procedure can be used only if you have at least two disk shelves for each node, which allows shelf-level auto-assignment of disks.

If you cannot use shelf-level auto-assignment, you must manually assign your remote disks so that each node has a remote pool of disks (pool 1).

The ONTAP automatic disk assignment feature assigns the disks on a shelf-by-shelf basis. For example:

- All the disks on site_B-shelf_2 are auto-assigned to pool1 of node_A_1
- All the disks on site_B-shelf_4 are auto-assigned to pool1 of node_A_2
- All the disks on site_A-shelf_2 are auto-assigned to pool1 of node_B_1
- All the disks on site_A-shelf_4 are auto-assigned to pool1 of node_B_2

You must "seed" the auto-assignment by specifying a single disk on each shelf.

Steps

1. From each node in the MetroCluster IP configuration, assign a remote disk to pool 1.

a. Display the list of unassigned disks:

```
disk show -host-adapter 0m -container-type unassigned
```

```
cluster_A::> disk show -host-adapter 0m -container-type unassigned
              Usable          Disk      Container  Container
Disk          Size Shelf Bay Type      Type      Name
Owner
-----
-----
6.23.0          -    23   0 SSD      unassigned -    -
6.23.1          -    23   1 SSD      unassigned -    -
.
.
.
node_A_2:0m.i1.2L51 -    21  14 SSD      unassigned -    -
node_A_2:0m.i1.2L64 -    21  10 SSD      unassigned -    -
.
.
.
48 entries were displayed.

cluster_A::>
```

b. Select a remote disk (0m) and assign ownership of the disk to pool 1 of the first node (for example, "node_A_1"):

```
disk assign -disk <disk_id> -pool 1 -owner <owner_node_name>
```

The disk-id must identify a disk on a remote shelf of owner_node_name.

The ONTAP disk auto-assignment feature assigns all disks on the remote shelf that contains the specified disk.

c. After waiting at least 60 seconds for disk auto-assignment to take place, verify that the remote disks on the shelf were auto-assigned to pool 1:

```
disk show -host-adapter 0m -container-type unassigned
```



The iSCSI connection used to access the remote disks appears as device 0m.

The following output shows that the disks on shelf 23 have now been assigned and no longer appear:

```

cluster_A::> disk show -host-adapter 0m -container-type unassigned
              Usable           Disk   Container   Container
Disk         Size Shelf Bay Type      Type        Name
Owner
-----
node_A_2:0m.i1.2L51      -    21  14 SSD      unassigned  -
node_A_2:0m.i1.2L64      -    21  10 SSD      unassigned  -
node_A_2:0m.i1.2L72      -    21  23 SSD      unassigned  -
node_A_2:0m.i1.2L74      -    21   1 SSD      unassigned  -
node_A_2:0m.i1.2L83      -    21  22 SSD      unassigned  -
node_A_2:0m.i1.2L90      -    21   7 SSD      unassigned  -
node_A_2:0m.i1.3L52      -    21   6 SSD      unassigned  -
node_A_2:0m.i1.3L59      -    21  13 SSD      unassigned  -
node_A_2:0m.i1.3L66      -    21  17 SSD      unassigned  -
node_A_2:0m.i1.3L73      -    21  12 SSD      unassigned  -
node_A_2:0m.i1.3L80      -    21   5 SSD      unassigned  -
node_A_2:0m.i1.3L81      -    21   2 SSD      unassigned  -
node_A_2:0m.i1.3L82      -    21  16 SSD      unassigned  -
node_A_2:0m.i1.3L91      -    21   3 SSD      unassigned  -
node_A_2:0m.i2.0L49      -    21  15 SSD      unassigned  -
node_A_2:0m.i2.0L50      -    21   4 SSD      unassigned  -
node_A_2:0m.i2.1L57      -    21  18 SSD      unassigned  -
node_A_2:0m.i2.1L58      -    21  11 SSD      unassigned  -
node_A_2:0m.i2.1L59      -    21  21 SSD      unassigned  -
node_A_2:0m.i2.1L65      -    21  20 SSD      unassigned  -
node_A_2:0m.i2.1L72      -    21   9 SSD      unassigned  -
node_A_2:0m.i2.1L80      -    21   0 SSD      unassigned  -
node_A_2:0m.i2.1L88      -    21   8 SSD      unassigned  -
node_A_2:0m.i2.1L90      -    21  19 SSD      unassigned  -
24 entries were displayed.

cluster_A::>

```

- d. Repeat these steps to assign pool 1 disks to the second node on site A (for example, "node_A_2").
- e. Repeat these steps on site B.

Enabling automatic drive assignment in ONTAP 9.4

About this task

In ONTAP 9.4, if you disabled automatic drive assignment as directed previously in this procedure, you must reenale it on all nodes.

[Considerations for automatic drive assignment and ADP systems in ONTAP 9.4 and later](#)

Steps

1. Enable automatic drive assignment:

```
storage disk option modify -node <node_name> -autoassign on
```

You must issue this command on all nodes in the MetroCluster IP configuration.

Mirroring the root aggregates

You must mirror the root aggregates to provide data protection.

About this task

By default, the root aggregate is created as RAID-DP type aggregate. You can change the root aggregate from RAID-DP to RAID4 type aggregate. The following command modifies the root aggregate for RAID4 type aggregate:

```
storage aggregate modify -aggregate <aggr_name> -raidtype raid4
```



On non-ADP systems, the RAID type of the aggregate can be modified from the default RAID-DP to RAID4 before or after the aggregate is mirrored.

Steps

1. Mirror the root aggregate:

```
storage aggregate mirror <aggr_name>
```

The following command mirrors the root aggregate for "controller_A_1":

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote MetroCluster site.

2. Repeat the previous step for each node in the MetroCluster configuration.

Related information

[Logical storage management](#)

Creating a mirrored data aggregate on each node

You must create a mirrored data aggregate on each node in the DR group.

About this task

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.
- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.

In systems using ADP, aggregates are created using partitions in which each drive is partitioned in to P1, P2 and P3 partitions.

- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.

Disk and aggregate management

- Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.

Steps

1. Display a list of available spares:

```
storage disk show -spare -owner <node_name>
```

2. Create the aggregate:

```
storage aggregate create -mirror true
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To ensure that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include



In the minimum supported configuration, in which a limited number of drives are available, you must use the `force-small-aggregate` option to allow the creation of a three disk RAID-DP aggregate.

- Checksum style to use for the aggregate
 - Type of drives to use
 - Size of drives to use
 - Drive speed to use
 - RAID type for RAID groups on the aggregate
 - Maximum number of drives that can be included in a RAID group
 - Whether drives with different RPM are allowed
- For more information about these options, see the `storage aggregate create` man page.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10
-node node_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_node_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate <aggregate-name>
```

Implementing the MetroCluster configuration

You must run the `metrocluster configure` command to start data protection in a MetroCluster configuration.

About this task

- There should be at least two non-root mirrored data aggregates on each cluster.

You can verify this with the `storage aggregate show` command.



If you want to use a single mirrored data aggregate, then see [Step 1](#) for instructions.

- The ha-config state of the controllers and chassis must be "mccip".

You issue the `metrocluster configure` command once on any of the nodes to enable the MetroCluster configuration. You do not need to issue the command on each of the sites or nodes, and it does not matter which node or site you choose to issue the command on.

The `metrocluster configure` command automatically pairs the two nodes with the lowest system IDs in each of the two clusters as disaster recovery (DR) partners. In a four-node MetroCluster configuration, there are two DR partner pairs. The second DR pair is created from the two nodes with higher system IDs.



You must **not** configure Onboard Key Manager (OKM) or external key management before you run the command `metrocluster configure`.

Steps

1. Configure the MetroCluster in the following format:

If your MetroCluster configuration has...	Then do this...
Multiple data aggregates	From any node's prompt, configure MetroCluster: <pre>metrocluster configure <node_name></pre>

A single mirrored data aggregate

- a. From any node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when you are prompted to continue into advanced mode and you see the advanced mode prompt (`*>`).

- b. Configure the MetroCluster with the `-allow-with-one-aggregate true` parameter:

```
metrocluster configure -allow-with-one-aggregate true <node_name>
```

- c. Return to the admin privilege level:

```
set -privilege admin
```



The best practice is to have multiple data aggregates. If the first DR group has only one aggregate and you want to add a DR group with one aggregate, you must move the metadata volume off the single data aggregate. For more information on this procedure, see [Moving a metadata volume in MetroCluster configurations](#).

The following command enables the MetroCluster configuration on all of the nodes in the DR group that contains "controller_A_1":

```
cluster_A::*> metrocluster configure -node-name controller_A_1  
  
[Job 121] Job succeeded: Configure is successful.
```

2. Verify the networking status on site A:

```
network port show
```

The following example shows the network port usage on a four-node MetroCluster configuration:

```
cluster_A::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper

controller_A_1						
	e0a	Cluster	Cluster	up	9000	auto/1000
	e0b	Cluster	Cluster	up	9000	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000
controller_A_2						
	e0a	Cluster	Cluster	up	9000	auto/1000
	e0b	Cluster	Cluster	up	9000	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000

```
14 entries were displayed.
```

3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Verify the configuration from site A:

```
metrocluster show
```

```
cluster_A::> metrocluster show
```

```
Configuration: IP fabric
```

Cluster	Entry Name	State

Local: cluster_A	Configuration state	configured
	Mode	normal
Remote: cluster_B	Configuration state	configured
	Mode	normal

b. Verify the configuration from site B:

```
metrocluster show
```

```
cluster_B::> metrocluster show
```

```
Configuration: IP fabric
```

Cluster	Entry Name	State
-----	-----	-----
Local: cluster_B	Configuration state	configured
	Mode	normal
Remote: cluster_A	Configuration state	configured
	Mode	normal

4. To avoid possible issues with nonvolatile memory mirroring, reboot each of the four nodes:

```
node reboot -node <node_name> -inhibit-takeover true
```

5. Issue the `metrocluster show` command on both clusters to again verify the configuration.

Configuring the second DR group in an eight-node configuration

Repeat the previous tasks to configure the nodes in the second DR group.

Creating unmirrored data aggregates

You can optionally create unmirrored data aggregates for data that does not require the redundant mirroring provided by MetroCluster configurations.

About this task

- Verify that you know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can verify that the correct drive type is selected.



In MetroCluster IP configurations, remote unmirrored aggregates are not accessible after a switchover



The unmirrored aggregates must be local to the node owning them.

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- *Disks and aggregates management* contains more information about mirroring aggregates.

Steps

1. Enable unmirrored aggregate deployment:

```
metrocluster modify -enable-unmirrored-aggr-deployment  
true
```

2. Verify that disk autoassignment is disabled:

```
disk option show
```

3. Install and cable the disk shelves that will contain the unmirrored aggregates.

You can use the procedures in the Installation and Setup documentation for your platform and disk shelves.

[ONTAP Hardware Systems Documentation](#)

4. Manually assign all disks on the new shelf to the appropriate node:

```
disk assign -disk <disk_id> -owner <owner_node_name>
```

5. Create the aggregate:

```
storage aggregate create
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To verify that the aggregate is created on a specific node, you should use the `-node` parameter or specify drives that are owned by that node.

You must also ensure that you are only including drives on the unmirrored shelf to the aggregate.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include
- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create` man page.

The following command creates a unmirrored aggregate with 10 disks:

```
controller_A_1::> storage aggregate create aggr1_controller_A_1
-diskcount 10 -node controller_A_1
[Job 15] Job is queued: Create aggr1_controller_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

6. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate <aggregate_name>
```

7. Disable unmirrored aggregate deployment:

```
metrocluster modify -enable-unmirrored-aggr-deployment false
```

8. Verify that disk autoassignment is enabled:

```
disk option show
```

Related information

[Disk and aggregate management](#)

Checking the MetroCluster configuration

You can check that the components and relationships in the MetroCluster configuration are working correctly.

About this task

You should do a check after initial configuration and after making any changes to the MetroCluster configuration.

You should also do a check before a negotiated (planned) switchover or a switchback operation.

If the `metrocluster check run` command is issued twice within a short time on either or both clusters, a conflict can occur and the command might not collect all data. Subsequent `metrocluster check show` commands do not show the expected output.

Steps

1. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```
cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results. To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"
```

```
cluster_A::> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

2. Display more detailed results from the most recent metrocluster check run command:

```
metrocluster check aggregate show
```

```
metrocluster check cluster show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```



The `metrocluster check show` commands show the results of the most recent `metrocluster check run` command. You should always run the `metrocluster check run` command prior to using the `metrocluster check show` commands so that the information displayed is current.

The following example shows the `metrocluster check aggregate show` command output for a healthy four-node MetroCluster configuration:

```
cluster_A::> metrocluster check aggregate show
```

Node	Aggregate	Check
Result		
-----	-----	-----
controller_A_1	controller_A_1_aggr0	mirroring-status
ok		disk-pool-allocation
ok		

```

ok                                     ownership-state
                                     controller_A_1_aggr1
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_1_aggr2
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_2_aggr0
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_2_aggr1
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_2_aggr2
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
18 entries were displayed.

```

The following example shows the `metrocluster check cluster show` command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

```
cluster_A::> metrocluster check cluster show
```

Cluster	Check	Result
mccint-fas9000-0102	negotiated-switchover-ready	not-applicable
	switchback-ready	not-applicable
	job-schedules	ok
	licenses	ok
	periodic-check-enabled	ok
mccint-fas9000-0304	negotiated-switchover-ready	not-applicable
	switchback-ready	not-applicable
	job-schedules	ok
	licenses	ok
	periodic-check-enabled	ok

10 entries were displayed.

Related information

[Disk and aggregate management](#)

[Network and LIF management](#)

Completing ONTAP configuration

After configuring, enabling, and checking the MetroCluster configuration, you can proceed to complete the cluster configuration by adding additional SVMs, network interfaces and other ONTAP functionality as needed.

Configure end-to-end encryption in a MetroCluster IP configuration

Beginning with ONTAP 9.15.1, you can configure end-to-end encryption on supported systems to encrypt back-end traffic, such as NVlog and storage replication data, between the sites in a MetroCluster IP configuration.

About this task

- You must be a cluster administrator to perform this task.
- Before you can configure end-to-end encryption, you must [Configure external key management](#).
- Review the supported systems and minimum ONTAP release required to configure end-to-end encryption in a MetroCluster IP configuration:

Minimum ONTAP release	Supported systems
ONTAP 9.17.1	<ul style="list-style-type: none"> • AFF A800, AFF C800 • AFF A20, AFF A30, AFF C30, AFF A50, AFF C60 • AFF A70, AFF A90, AFF A1K, AFF C80 • FAS50, FAS70, FAS90
ONTAP 9.15.1	<ul style="list-style-type: none"> • AFF A400 • AFF C400 • FAS8300 • FAS8700

Enable end-to-end encryption

Perform the following steps to enable end-to-end encryption.

Steps

1. Verify the health of the MetroCluster configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- b. After the `metrocluster check run` operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

- c. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

- d. Verify that there are no health alerts:

```
system health alert show
```

2. Verify that external key management is configured on both clusters:

```
security key-manager external show-status
```

3. Enable end-to-end encryption for each DR group:

```
metrocluster modify -is-encryption-enabled true -dr-group-id  
<dr_group_id>
```

Example

```
cluster_A:::*> metrocluster modify -is-encryption-enabled true -dr-group  
-id 1  
Warning: Enabling encryption for a DR Group will secure NVLog and  
Storage  
        replication data sent between MetroCluster nodes and have an  
impact on  
        performance. Do you want to continue? {y|n}: y  
[Job 244] Job succeeded: Modify is successful.
```

Repeat this step for each DR group in the configuration.

4. Verify that end-to-end encryption is enabled:

```
metrocluster node show -fields is-encryption-enabled
```

Example

```
cluster_A::*> metrocluster node show -fields is-encryption-enabled

dr-group-id cluster      node      configuration-state is-encryption-
enabled
-----
1           cluster_A    node_A_1  configured         true
1           cluster_A    node_A_2  configured         true
1           cluster_B    node_B_1  configured         true
1           cluster_B    node_B_2  configured         true
4 entries were displayed.
```

Disable end-to-end encryption

Perform the following steps to disable end-to-end encryption.

Steps

1. Verify the health of the MetroCluster configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- b. After the metrocluster check run operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

- c. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

- d. Verify that there are no health alerts:

```
system health alert show
```

2. Verify that external key management is configured on both clusters:

```
security key-manager external show-status
```

3. Disable end-to-end encryption on each DR group:

```
metrocluster modify -is-encryption-enabled false -dr-group-id  
<dr_group_id>
```

Example

```
cluster_A:::*> metrocluster modify -is-encryption-enabled false -dr-group  
-id 1  
[Job 244] Job succeeded: Modify is successful.
```

Repeat this step for each DR group in the configuration.

4. Verify that end-to-end encryption is disabled:

```
metrocluster node show -fields is-encryption-enabled
```

Example

```
cluster_A::*> metrocluster node show -fields is-encryption-enabled

dr-group-id cluster      node      configuration-state is-encryption-
enabled
-----
1           cluster_A    node_A_1  configured         false
1           cluster_A    node_A_2  configured         false
1           cluster_B    node_B_1  configured         false
1           cluster_B    node_B_2  configured         false
4 entries were displayed.
```

Set up MetroCluster Tiebreaker or ONTAP Mediator for a MetroCluster IP configuration

You can download and install on a third site either the MetroCluster Tiebreaker software, or, beginning with ONTAP 9.7, the ONTAP Mediator.

Before you begin

You must have a Linux host available that has network connectivity to both clusters in the MetroCluster configuration. The specific requirements are in the MetroCluster Tiebreaker or ONTAP Mediator documentation.

If you are connecting to an existing Tiebreaker or ONTAP Mediator instance, you need the username, password, and IP address of the Tiebreaker or Mediator.

If you must install a new instance of the ONTAP Mediator, follow the directions to install and configure the software.

[Configure ONTAP Mediator for unplanned automatic switchover](#)

If you must install a new instance of the Tiebreaker software, follow the [directions to install and configure the software](#).

About this task

You cannot use both the MetroCluster Tiebreaker software and the ONTAP Mediator with the same MetroCluster configuration.

[Considerations for using ONTAP Mediator or MetroCluster Tiebreaker](#)

Step

1. Configure ONTAP Mediator or the Tiebreaker software:
 - If you are using an existing instance of the ONTAP Mediator, add ONTAP Mediator to ONTAP:

```
metrocluster configuration-settings mediator add -mediator-address ip-
```

address-of-mediator-host

- If you are using the Tiebreaker software, refer to the [Tiebreaker documentation](#).

Backup cluster configuration files in a MetroCluster IP configuration

You can provide additional protection for the cluster configuration backup files by specifying a remote URL (either HTTP or FTP) where the configuration backup files will be uploaded in addition to the default locations in the local cluster.

Step

1. Set the URL of the remote destination for the configuration backup files:

```
system configuration backup settings modify URL-of-destination
```

The [Cluster Management with the CLI](#) contains additional information under the section *Managing configuration backups*.

Configure the MetroCluster software using System Manager

Set up a MetroCluster IP site with ONTAP System Manager

Beginning with ONTAP 9.8, you can use System Manager to set up a MetroCluster IP site.

A MetroCluster site consists of two clusters. Typically, the clusters are located in different geographical locations.

Before you begin

- Your system should already be installed and cabled according to the [Installation and Setup Instructions](#) that came with the system.
- Cluster network interfaces should be configured on each node of each cluster for intra-cluster communication.

Assign a node-management IP address

Windows System

You should connect your Windows computer to the same subnet as the controllers. This automatically assigns a node-management IP address to your system.

Steps

1. From the Windows system, open the **Network** drive to discover the nodes.
2. Double-click the node to launch the cluster setup wizard.

Other systems

You should configure the node-management IP address for one of the nodes in your cluster. You can use this node-management IP address to launch the cluster set up wizard.

See [Creating the cluster on the first node](#) for information about assigning a node-management IP address.

Initialize and configure the cluster

You initialize the cluster by setting an administrative password for the cluster and setting up the cluster management and node management networks. You can also configure services like a domain name server (DNS) to resolve host names and an NTP server to synchronize time.

Steps

1. On a web browser, enter the node-management IP address that you have configured: "https://node-management-IP"

System Manager automatically discovers the remaining nodes in the cluster.

2. In the **Initialize Storage System** window, perform the following:
 - a. Enter cluster management network configuration data.
 - b. Enter Node management IP addresses for all the nodes.
 - c. Provide DNS details.
 - d. In the **Other** section, select the check box labeled **Use time service (NTP)** to add the time servers.

When you click **Submit**, wait for the cluster to be created and configured. Then, a validation process occurs.

What's Next?

After both clusters have been set up, initialized, and configured, perform the [Set up MetroCluster IP peering](#) procedure.

Configure ONTAP on a new cluster video



Set up MetroCluster IP peering with ONTAP System Manager

Beginning with ONTAP 9.8, you can manage MetroCluster IP configuration operations

with System Manager. After setting up two clusters, you set up peering between them.

Before you begin

Set up two clusters. See the [Set up a MetroCluster IP site](#) procedure.

Certain steps of this process are performed by different system administrators located at the geographical sites of each cluster. For the purposes of explaining this process, the clusters are called "Site A cluster" and "Site B cluster".

Perform the peering process from Site A

This process is performed by a system administrator at Site A.

Steps

1. Log in to Site A cluster.
2. In System Manager, select **Dashboard** from the left navigation column to display the cluster overview.

The dashboard shows the details for this cluster (Site A). In the **MetroCluster** section, Site A cluster is shown on the left.

3. Click **Attach Partner Cluster**.
4. Enter the details of the network interfaces that allow the nodes in Site A cluster to communicate with the nodes in Site B cluster.
5. Click **Save and Continue**.
6. On the **Attach Partner Cluster** window, select **I do not have a passphrase**. This lets you generate a passphrase.
7. Copy the generated passphrase and share it with the system administrator at Site B.
8. Select **Close**.

Perform the peering process from Site B

This process is performed by a system administrator at Site B.

Steps

1. Log in to Site B cluster.
2. In System Manager, select **Dashboard** to display the cluster overview.

The dashboard shows the details for this cluster (Site B). In the MetroCluster section, Site B cluster is shown on the left.

3. Click **Attach Partner Cluster** to start the peering process.
4. Enter the details of the network interfaces that allow the nodes in Site B cluster to communicate with the nodes in Site A cluster.
5. Click **Save and Continue**.
6. On the **Attach Partner Cluster** window, select **I have a passphrase**. This lets you enter the passphrase that you received from the system administrator at Site A.
7. Select **Peer** to complete the peering process.

What's next?

After the peering process successfully completes, you configure the clusters. See [Configure a MetroCluster IP site](#).

Configure a MetroCluster IP site with ONTAP System Manager

Beginning with ONTAP 9.8, you can manage MetroCluster IP configuration operations with System Manager. This involves setting up two clusters, performing cluster peering, and configuring the clusters.

Before you begin

Complete the following procedures:

- [Set up a MetroCluster IP site](#)
- [Set up MetroCluster IP peering](#)

Configure the connection between clusters

Steps

1. Log in to System Manager on one of the sites, and select **Dashboard**.

In the **MetroCluster** section, the graphic shows the two clusters that you set up and peered for the MetroCluster sites. The cluster you are working from (local cluster) is shown on the left.

2. Click **Configure MetroCluster**. From this window, perform the following steps:
 - a. The nodes for each cluster in the MetroCluster configuration are shown. Use the drop-down lists to select the nodes in the local cluster that will be disaster recovery partners with the nodes in the remote cluster.
 - b. Click the check box if you want to configure ONTAP Mediator. See [Configure ONTAP Mediator](#).
 - c. If both clusters have a license to enable encryption, the **Encryption** section is displayed.

To enable encryption, enter a passphrase.

- d. Click the check box if you want to configure MetroCluster with a shared layer 3 network.



The HA partner nodes and network switches connecting to the nodes must have a matching configuration.

3. Click **Save** to configure the MetroCluster sites.

On the **Dashboard**, in the **MetroCluster** section, the graphic shows a check mark on the link between the two clusters, indicating a healthy connection.

Configure ONTAP Mediator for unplanned automatic switchover

ONTAP Mediator installation requirements for MetroCluster IP configurations

Your environment must meet certain requirements.

The following requirements apply to one disaster recovery group (DR group). Learn more about [DR groups](#).

- If you plan on updating your Linux version, do so before you install the most current version of ONTAP Mediator.
- The ONTAP Mediator and MetroCluster Tiebreaker software should not both be used with the same MetroCluster configuration.
- ONTAP Mediator must be installed on a Linux host at a separate location from the MetroCluster sites.

The connectivity between the ONTAP Mediator and each site must be two separate failure domains.

- Automatic unplanned switchover is supported in ONTAP 9.7 and later.
- Beginning with ONTAP 9.18.1 and ONTAP Mediator 1.11, a single ONTAP Mediator instance can manage up to ten MetroCluster configurations simultaneously. In earlier releases, ONTAP Mediator can support up to five MetroCluster configurations simultaneously.
- Beginning with ONTAP 9.18.1, IPv6 is supported for ONTAP Mediator 1.11 or later in a MetroCluster IP configuration.

Network requirements for using ONTAP Mediator in a MetroCluster configuration

To install ONTAP Mediator in a MetroCluster configuration, you must make sure that the configuration meets several network requirements.

- Latency

Maximum latency of less than 75ms (RTT).

Jitter must be no more than 5ms.

- MTU

The MTU size must be at least 1400.

- Packet loss

For both Internet Control Message Protocol (ICMP) and TCP traffic, packet loss must be less than 0.01%.

- Bandwidth

The link between ONTAP Mediator and one DR group must have at least 20Mbps of bandwidth.

- Independent connectivity

Independent connectivity between each site and the ONTAP Mediator is required. A failure in one site must not interrupt the IP connectivity between the other two unaffected sites.

Host requirements for ONTAP Mediator in a MetroCluster configuration

You must ensure that the configuration meets several host requirements.

- ONTAP Mediator must be installed at an external site that is physically separated from the two ONTAP clusters.
- ONTAP Mediator does not require more than the host operating system's minimum requirements for CPU and memory (RAM).

- In addition to the host operating system's minimum requirements, at least 30GB of additional usable disk space must be available.
 - Each DR group requires up to 200MB of disk space.

Firewall requirements for ONTAP Mediator

ONTAP Mediator uses a number of ports to communicate with specific services.

If you are using a third-party firewall:

- HTTPS access must be enabled.
- It must be configured to allow access on ports 31784 and 3260.

When using the default Red Hat or CentOS firewall, the firewall is automatically configured during Mediator installation.

The following table lists the ports that you must allow in your firewall:



- The iSCSI port is only required in a MetroCluster IP configuration.
- The 22/tcp port is not required for normal operation but you can enable it temporarily for maintenance and disable it when the maintenance session has finished.

Port/services	Source	Direction	Destination	Purpose
22/tcp	Management host	Inbound	ONTAP Mediator	SSH / ONTAP Mediator management
31784/tcp	cluster-mgmt and node-mgmt LIFs	Inbound	ONTAP Mediator web server	REST API (HTTPS)
3260/tcp	node-mgmt LIFs	Inbound	ONTAP Mediator iSCSI targets	iSCSI data connection for mailboxes

Guidelines for upgrading ONTAP Mediator in a MetroCluster configuration

If you are upgrading ONTAP Mediator you must meet the Linux version requirements and follow guidelines for the upgrade.

- ONTAP Mediator can be upgraded from version from an immediately prior version to the current version.
- All Mediator versions are supported on MetroCluster IP configurations running ONTAP 9.7 or later.

[Install or upgrade ONTAP Mediator](#)

After the upgrade

After the Mediator and operating system upgrade is complete, you should issue the `storage iscsi-initiator show` command to confirm that the Mediator connections are up.

Set up the ONTAP Mediator for a MetroCluster IP configuration

You must configure the ONTAP Mediator on the ONTAP node to use it in a MetroCluster IP configuration.

Before you begin

- ONTAP Mediator must have been successfully installed on a network location that can be reached by both MetroCluster sites.

[Install or upgrade ONTAP Mediator](#)

- You must have the IP address of the host running ONTAP Mediator.
- You must have the username and password for ONTAP Mediator.
- All nodes of the MetroCluster IP configuration must be online.



Beginning with ONTAP 9.12.1, you can enable the MetroCluster automatic forced switchover feature in a MetroCluster IP configuration. This feature is an extension of the Mediator-assisted unplanned switchover. Before you enable this feature, review the [Risks and limitations of using MetroCluster automatic forced switchover](#).

About this task

- This task enables automatic unplanned switchover by default.
- This task can be performed on the ONTAP interface of any node in the MetroCluster IP configuration.
- Beginning with ONTAP 9.18.1 and ONTAP Mediator 1.11, a single ONTAP Mediator instance can manage up to ten MetroCluster configurations simultaneously. In earlier releases, ONTAP Mediator can support up to five MetroCluster configurations simultaneously.

Steps

1. Add ONTAP Mediator to ONTAP. The steps depend on whether you want to use an IPv4 or IPv6 address.



- You must be running ONTAP 9.18.1 or later and ONTAP Mediator 1.11 or later to use IPv6.
- If you enable IPv6 on a cluster, you cannot disable it later.

Use IPv4

- a. Run the following command to add the ONTAP Mediator:

```
metrocluster configuration-settings mediator add -mediator-address  
<mediator_host_ip_address>
```



You are prompted for the username and password for the Mediator admin user account.

Use IPv6

- a. Run the following command on both clusters:

```
network options ipv6 modify -enabled true
```

- b. Configure the node-mgmt IP address with IPv6 addresses on all four nodes.
- c. Add the ONTAP Mediator:

```
metrocluster configuration-settings mediator add -mediator-address  
<mediator_host_ipv6_ip_address>
```



You are prompted for the username and password for the Mediator admin user account.

2. Verify that the automatic switchover feature is enabled:

```
metrocluster show
```

3. Verify that the Mediator is now running.

- a. Show the Mediator virtual disks:

```
storage disk show -container-type mediator
```

```
cluster_A::> storage disk show -container-type mediator
                Usable          Disk      Container
Container
Disk           Size Shelf Bay Type      Type      Name
Owner
-----
NET-1.5        -      -   - VMDISK  mediator  -
node_A_2
NET-1.6        -      -   - VMDISK  mediator  -
node_B_1
NET-1.7        -      -   - VMDISK  mediator  -
node_B_2
NET-1.8        -      -   - VMDISK  mediator  -
node_A_1
```

b. Set the privilege mode to advanced:

```
set advanced
```

```
cluster_A::> set advanced
```

c. Display the initiators labelled as mediator:

```
storage iscsi-initiator show -label mediator
```

```

cluster_A::*> storage iscsi-initiator show -label mediator
(storage iscsi-initiator show)
+
Status
Node Type Label      Target Portal      Target Name
Admin/Op
-----
node_A_1
  mailbox
    mediator 1.1.1.1      iqn.2012-
05.local:mailbox.target.6616cd3f-9ef1-11e9-aada-
00a098ccf5d8:a05e1ffb-9ef1-11e9-8f68- 00a098cbca9e:1 up/up
node_A_2
  mailbox
    mediator 1.1.1.1      iqn.2012-
05.local:mailbox.target.6616cd3f-9ef1-11e9-aada-
00a098ccf5d8:a05e1ffb-9ef1-11e9-8f68-00a098cbca9e:1 up/up

```

d. Verify the state of the automatic unplanned switchover (AUSO) failure domain:

```
metrocluster show
```



The following example output applies to ONTAP 9.13.1 and later. For ONTAP 9.12.1 and earlier, the AUSO failure domain state should be `auso-on-cluster-disaster`.

```

cluster_A::> metrocluster show
Cluster              Entry Name          State
-----
Local: cluster_A    Configuration state configured
Mode                 normal
AUSO Failure Domain auso-on-dr-group-
disaster
Remote: cluster_B   Configuration state configured
Mode                 normal
AUSO Failure Domain auso-on-dr-group-
disaster

```

4. Optionally, configure MetroCluster automatic forced switchover.

You can only use the following command in advanced privilege level.



Before using this command, review the [Risks and limitations of using MetroCluster automatic forced switchover](#).

```
metrocluster modify -allow-auto-forced-switchover true
```

Example

```
cluster_A::*> metrocluster modify -allow-auto-forced-switchover true
```

Remove the ONTAP Mediator from a MetroCluster IP configuration

You can unconfigure ONTAP Mediator from the MetroCluster IP configuration.

Before you begin

You must have successfully installed and configured ONTAP Mediator on a network location that can be reached by both MetroCluster sites.

Steps

1. Unconfigure ONTAP Mediator by using the following command:

```
metrocluster configuration-settings mediator remove
```

You are prompted for the user name and password for the ONTAP Mediator admin user account.



If the ONTAP Mediator is down, the `metrocluster configuration-settings mediator remove` command still prompts you to enter the user name and password for the ONTAP Mediator admin user account and removes ONTAP Mediator from the MetroCluster configuration.

- a. Check if there are any broken disks by using the following command:

```
disk show -broken
```

Example

```
There are no entries matching your query.
```

2. Confirm that ONTAP Mediator has been removed from the MetroCluster configuration by running the following commands on both clusters:

- a. `metrocluster configuration-settings mediator show`

Example

```
This table is currently empty.
```

- b. `storage iscsi-initiator show -label mediator`

Example

There are no entries matching your query.

Connect a MetroCluster IP configuration to a different ONTAP Mediator instance

If you want to connect the MetroCluster nodes to a different ONTAP Mediator instance, you must unconfigure and then reconfigure the Mediator connection in the ONTAP software.

Before you begin

You need the username, password, and IP address of the new ONTAP Mediator instance.

About this task

These commands can be issued from any node in the MetroCluster configuration.

Steps

1. Remove the current ONTAP Mediator from the MetroCluster configuration:

```
metrocluster configuration-settings mediator remove
```

2. Establish the new ONTAP Mediator connection to the MetroCluster configuration:

```
metrocluster configuration-settings mediator add -mediator-address ip-address-of-mediator-host
```

How the ONTAP Mediator supports automatic unplanned switchover in MetroCluster IP configurations

ONTAP Mediator provides mailbox LUNs to store state information about the MetroCluster IP nodes. These LUNs are co-located with ONTAP Mediator, which runs on a Linux host physically separate from the MetroCluster sites. The MetroCluster IP nodes can use the mailbox information to monitor the state of their disaster recovery (DR) partners and implement a Mediator-assisted unplanned switchover (MAUSO) in the case of a disaster.



MAUSO is not supported in MetroCluster FC configurations.

When a node detects a site failure requiring a switchover, it takes steps to confirm that the switchover is appropriate and, if so, performs the switchover. By default, a MAUSO is initiated for the following scenarios:

- Both SyncMirror mirroring and DR mirroring of each node's nonvolatile cache is operating and the caches and mirrors are synchronized at the time of the failure.
- None of the nodes at the surviving site are in takeover state.
- If a site disaster occurs. A site disaster is a failure of *all* nodes at the same site.

A MAUSO is *not* initiated in the following shutdown scenarios:

- You initiate a shutdown. For example, when you:

- Halt the nodes
- Reboot the nodes

Learn about the MAUSO features available with each ONTAP 9 release.

Beginning with...	Description
ONTAP 9.13.1	<ul style="list-style-type: none"> • A MAUSO is initiated if a default scenario occurs and a fan or hardware failure initiates an environmental shutdown. Examples of hardware failures include a high or low temperature, or a power supply unit, NVRAM battery, or Service Processor heartbeat failure. • The default value for the failure domain is set to "auso-on-dr-group" in a MetroCluster IP configuration. For ONTAP 9.12.1 and earlier, the default value is set to "auso-on-cluster-disaster". <p>In an eight-node MetroCluster IP configuration, "auso-on-dr-group" triggers a MAUSO either on failure of the cluster or a HA pair in one DR group. For a HA pair, both nodes must fail at the same time.</p> <p>Optionally, you can change the failure domain setting to the "auso-on-cluster-disaster" domain using the <code>metrocluster modify -auto-switchover-failure-domain auso-on-cluster-disaster</code> command to trigger a MAUSO only if there are HA node pair failures in both DR groups.</p> <ul style="list-style-type: none"> • You can change the behavior to force a MAUSO even if NVRAM is not in sync at the time of the failure.
ONTAP 9.12.1	<p>You can enable the MetroCluster automatic forced switchover feature in a MetroCluster IP configuration by using the <code>metrocluster modify -allow-auto-forced -switchover true</code> command.</p> <p>Switchover upon detection of a site failure happens automatically when you enable the MetroCluster automatic forced switchover feature. You can use this feature to supplement the MetroCluster IP automatic switchover capability.</p> <p>Risks and limitations of using MetroCluster automatic forced switchover</p> <p>When you allow a MetroCluster IP configuration to operate in automatic forced switchover mode, the following known issue might lead to data loss:</p> <ul style="list-style-type: none"> • The nonvolatile memory in the storage controllers is not mirrored to the remote DR partner on the partner site. <p>Caution: You might encounter scenarios that are not mentioned. NetApp is not responsible for any data corruption, data loss, or other damage that might arise from enabling the MetroCluster automatic forced switchover feature. Do not use the MetroCluster automatic forced switchover feature if the risks and limitations are not acceptable to you.</p>

Manage the ONTAP Mediator with System Manager in MetroCluster IP configurations

Using System Manager, you can perform tasks to manage ONTAP Mediator.

About these tasks

Beginning with ONTAP 9.8, you can use System Manager as a simplified interface for managing a four-node MetroCluster IP configuration, which can include an ONTAP Mediator installed in a third location.

Beginning with ONTAP 9.14.1, you can use System Manager to also perform these operations for an eight-node MetroCluster IP site. Although you can't set up or expand an eight-node system with System Manager, if you have already set up an eight-node MetroCluster IP system, then you can perform these operations.

Perform the following tasks to manage ONTAP Mediator.

To perform this task...	Take these actions...
Configure ONTAP Mediator	<p>Both clusters at the MetroCluster sites should be up and peered.</p> <p>Steps</p> <ol style="list-style-type: none">1. In System Manager in ONTAP 9.8, select Cluster > Settings.2. In the Mediator section, click the .3. On the Configure Mediator window, click Add+.4. Enter the configuration details for ONTAP Mediator. <p>You can enter the following details while configuring ONTAP Mediator with System Manager.</p> <ul style="list-style-type: none">◦ The IP address of ONTAP Mediator.◦ The user name.◦ The password.
Enable or disable Mediator-assisted Automatic Switchover (MAUSO)	<p>Steps</p> <ol style="list-style-type: none">1. In System Manager, click Dashboard.2. Scroll to the MetroCluster section.3. Click  next to the MetroCluster site name.4. Select Enable or Disable.5. Enter the administrator user name and password, then click Enable or Disable. <div data-bbox="625 1640 1463 1787"><p>You can enable or disable ONTAP Mediator when it can be reached and both sites are in "Normal" mode. ONTAP Mediator is still reachable when MAUSO is enabled or disabled if the MetroCluster system is healthy.</p></div>

Remove ONTAP Mediator from the MetroCluster configuration	<p>Steps</p> <ol style="list-style-type: none"> 1. In System Manager, click Dashboard. 2. Scroll to the MetroCluster section. 3. Click  next to the MetroCluster site name. 4. Select Remove Mediator. 5. Enter the administrator user name and password, then click Remove.
Check the health of ONTAP Mediator	Perform the System Manager specific steps in Verify the health of a MetroCluster configuration .
Perform a switchover and a switchback	Perform the steps in Use System Manger to perform switchover and switchback (MetroCluster IP configurations only) .

Test the ONTAP node switchover for your MetroCluster IP configuration

You can test failure scenarios to confirm the correct operation of the MetroCluster configuration.

Verifying negotiated switchover

You can test the negotiated (planned) switchover operation to confirm uninterrupted data availability.

About this task

This test validates that data availability is not affected (except for SMB and Fibre Channel protocols) by switching the cluster over to the second data center.

This test should take about 30 minutes.

This procedure has the following expected results:

- The `metrocluster switchover` command will present a warning prompt.

If you respond `yes` to the prompt, the site the command is issued from will switch over the partner site.

For MetroCluster IP configurations:

- For ONTAP 9.4 and earlier:
 - Mirrored aggregates will become degraded after the negotiated switchover.
- For ONTAP 9.5 and later:
 - Mirrored aggregates will remain in normal state if the remote storage is accessible.
 - Mirrored aggregates will become degraded after the negotiated switchover if access to the remote storage is lost.
- For ONTAP 9.8 and later:
 - Unmirrored aggregates that are located at the disaster site will become unavailable if access to the

remote storage is lost. This might lead to a controller outage.

Steps

1. Confirm that all nodes are in the configured state and normal mode:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
```

Cluster	Configuration State	Mode
-----	-----	

Local: cluster_A	configured	normal
Remote: cluster_B	configured	normal

2. Begin the switchover operation:

```
metrocluster switchover
```

```
cluster_A::> metrocluster switchover
Warning: negotiated switchover is about to start. It will stop all the
data Vservers on cluster "cluster_B" and
automatically re-start them on cluster "cluster_A". It will finally
gracefully shutdown cluster "cluster_B".
```

3. Confirm that the local cluster is in the configured state and switchover mode:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
```

Cluster	Configuration State	Mode
-----	-----	

Local: cluster_A	configured	switchover
Remote: cluster_B	not-reachable	-
configured	normal	

4. Confirm that the switchover operation was successful:

```
metrocluster operation show
```

```
cluster_A::> metrocluster operation show
  Operation: switchover
    State: successful
  Start Time: 2/6/2016 13:28:50
  End Time: 2/6/2016 13:29:41
  Errors: -
```

5. Use the `vserver show` and `network interface show` commands to verify that DR SVMs and LIFs have come online.

Verifying healing and manual switchback

You can test the healing and manual switchback operations to verify that data availability is not affected (except for SMB and Solaris FC configurations) by switching back the cluster to the original data center after a negotiated switchover.

About this task

This test should take about 30 minutes.

The expected result of this procedure is that services should be switched back to their home nodes.

The healing steps are not required on systems running ONTAP 9.5 or later, on which healing is performed automatically after a negotiated switchover. On systems running ONTAP 9.6 and later, healing is also performed automatically after unscheduled switchover.

Steps

1. If the system is running ONTAP 9.4 or earlier, heal the data aggregate:

```
metrocluster heal aggregates
```

The following example shows the successful completion of the command:

```
cluster_A::> metrocluster heal aggregates
[Job 936] Job succeeded: Heal Aggregates is successful.
```

2. If the system is running ONTAP 9.4 or earlier, heal the root aggregate:

```
metrocluster heal root-aggregates
```

This step is required on the following configurations:

- MetroCluster FC configurations.
 - MetroCluster IP configurations running ONTAP 9.4 or earlier.
- The following example shows the successful completion of the command:

```
cluster_A::> metrocluster heal root-aggregates
[Job 937] Job succeeded: Heal Root Aggregates is successful.
```

3. Verify that healing is completed:

```
metrocluster node show
```

The following example shows the successful completion of the command:

```
cluster_A::> metrocluster node show
DR                               Configuration  DR
Group Cluster Node              State          Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured    enabled    heal roots
completed
      cluster_B
      node_B_2      unreachable  -          switched over
42 entries were displayed.
```

If the automatic healing operation fails for any reason, you must issue the `metrocluster heal` commands manually as done in ONTAP versions prior to ONTAP 9.5. You can use the `metrocluster operation show` and `metrocluster operation history show -instance` commands to monitor the status of healing and determine the cause of a failure.

4. Verify that all aggregates are mirrored:

```
storage aggregate show
```

The following example shows that all aggregates have a RAID Status of mirrored:

```

cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate Size      Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
data_cluster
      4.19TB      4.13TB   2% online    8 node_A_1  raid_dp,
mirrored,
normal

root_cluster
      715.5GB    212.7GB  70% online    1 node_A_1  raid4,
mirrored,
normal

cluster_B Switched Over Aggregates:
Aggregate Size      Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
data_cluster_B
      4.19TB      4.11TB   2% online    5 node_A_1  raid_dp,
mirrored,
normal

root_cluster_B    -          -      - unknown    - node_A_1  -

```

5. Check the status of switchback recovery:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR                               Configuration  DR
Group Cluster Node              State          Mirroring Mode
-----
-----
1      cluster_A
      node_A_1          configured    enabled    heal roots
completed
      cluster_B
      node_B_2          configured    enabled    waiting for
switchback                                     recovery

2 entries were displayed.

```

6. Perform the switchback:

```
metrocluster switchback
```

```
cluster_A::> metrocluster switchback  
[Job 938] Job succeeded: Switchback is successful. Verify switchback
```

7. Confirm status of the nodes:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show  
DR                               Configuration  DR  
Group Cluster Node              State          Mirroring Mode  
-----  
1      cluster_A  
        node_A_1      configured    enabled      normal  
        cluster_B  
        node_B_2      configured    enabled      normal  
  
2 entries were displayed.
```

8. Confirm status of the MetroCluster operation:

```
metrocluster operation show
```

The output should show a successful state.

```
cluster_A::> metrocluster operation show  
Operation: switchback  
State: successful  
Start Time: 2/6/2016 13:54:25  
End Time: 2/6/2016 13:56:15  
Errors: -
```

Verifying operation after power line disruption

You can test the MetroCluster configuration's response to the failure of a PDU.

About this task

The best practice is for each power supply unit (PSU) in a component to be connected to separate power supplies. If both PSUs are connected to the same power distribution unit (PDU) and an electrical disruption occurs, the site could down or a complete shelf might become unavailable. Failure of one power line is tested to confirm that there is no cabling mismatch that could cause a service disruption.

This test should take about 15 minutes.

This test requires turning off power to all left-hand PDUs and then all right-hand PDUs on all of the racks containing the MetroCluster components.

This procedure has the following expected results:

- Errors should be generated as the PDUs are disconnected.
- No failover or loss of service should occur.

Steps

1. Turn off the power of the PDUs on the left-hand side of the rack containing the MetroCluster components.
2. Monitor the result on the console:

```
system environment sensors show -state fault
```

```
storage shelf show -errors
```

```
cluster_A::> system environment sensors show -state fault

Node Sensor                State Value/Units Crit-Low Warn-Low Warn-Hi
Crit-Hi
-----
node_A_1
  PSU1                    fault
                        PSU_OFF
  PSU1 Pwr In OK          fault
                        FAULT
node_A_2
  PSU1                    fault
                        PSU_OFF
  PSU1 Pwr In OK          fault
                        FAULT

4 entries were displayed.

cluster_A::> storage shelf show -errors
  Shelf Name: 1.1
  Shelf UID: 50:0a:09:80:03:6c:44:d5
  Serial Number: SHFHU1443000059

Error Type                Description
-----
Power                    Critical condition is detected in storage shelf
power supply unit "1". The unit might fail.Reconnect PSU1
```

3. Turn the power back on to the left-hand PDUs.
4. Make sure that ONTAP clears the error condition.

5. Repeat the previous steps with the right-hand PDUs.

Verifying operation after loss of a single storage shelf

You can test the failure of a single storage shelf to verify that there is no single point of failure.

About this task

This procedure has the following expected results:

- An error message should be reported by the monitoring software.
- No failover or loss of service should occur.
- Mirror resynchronization starts automatically after the hardware failure is restored.

Steps

1. Check the storage failover status:

```
storage failover show
```

```
cluster_A::> storage failover show

Node           Partner           Possible State Description
-----
node_A_1       node_A_2          true      Connected to node_A_2
node_A_2       node_A_1          true      Connected to node_A_1
2 entries were displayed.
```

2. Check the aggregate status:

```
storage aggregate show
```

```
cluster_A::> storage aggregate show
```

```
cluster Aggregates:
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID

node_A_1data01_mirrored	4.15TB	3.40TB	18%	online	3	node_A_1	
raid_dp,							
mirrored,							
normal							
node_A_1root	707.7GB	34.29GB	95%	online	1	node_A_1	
raid_dp,							
mirrored,							
normal							
node_A_2_data01_mirrored	4.15TB	4.12TB	1%	online	2	node_A_2	
raid_dp,							
mirrored,							
normal							
node_A_2_data02_unmirrored	2.18TB	2.18TB	0%	online	1	node_A_2	
raid_dp,							
normal							
node_A_2_root	707.7GB	34.27GB	95%	online	1	node_A_2	
raid_dp,							
mirrored,							
normal							

3. Verify that all data SVMs and data volumes are online and serving data:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0,!MDV*
```

```
cluster_A::> vserver show -type data
Vserver      Type      Subtype      Admin      Operational  Root
Aggregate
-----
SVM1         data      sync-source      running      SVM1_root
node_A_1_data01_mirrored
SVM2         data      sync-source      running      SVM2_root
node_A_2_data01_mirrored
```

```
cluster_A::> network interface show -fields is-home false
There are no entries matching your query.
```

```
cluster_A::> volume show !vol0,!MDV*
```

```
Vserver      Volume      Aggregate      State      Type      Size
Available Used%
-----
SVM1
          SVM1_root
                    node_A_1data01_mirrored
                    online      RW      10GB
9.50GB      5%
SVM1
          SVM1_data_vol
                    node_A_1data01_mirrored
                    online      RW      10GB
9.49GB      5%
SVM2
          SVM2_root
                    node_A_2_data01_mirrored
                    online      RW      10GB
9.49GB      5%
SVM2
          SVM2_data_vol
                    node_A_2_data02_unmirrored
                    online      RW      1GB
972.6MB      5%
```

4. Identify a shelf in Pool 1 for node "node_A_2" to power off to simulate a sudden hardware failure:

```
storage aggregate show -r -node node-name !*root
```

The shelf you select must contain drives that are part of a mirrored data aggregate.

In the following example, shelf ID "31" is selected to fail.

```
cluster_A::> storage aggregate show -r -node node_A_2 !*root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirrored) (block
checksums)
Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block
checksums)

Physical
Position Disk Pool Type RPM Usable
Size Status Size
-----
-----
dparity 2.30.3 0 BSAS 7200 827.7GB
828.0GB (normal)
parity 2.30.4 0 BSAS 7200 827.7GB
828.0GB (normal)
data 2.30.6 0 BSAS 7200 827.7GB
828.0GB (normal)
data 2.30.8 0 BSAS 7200 827.7GB
828.0GB (normal)
data 2.30.5 0 BSAS 7200 827.7GB
828.0GB (normal)

Plex: /node_A_2_data01_mirrored/plex4 (online, normal, active, pool1)
RAID Group /node_A_2_data01_mirrored/plex4/rg0 (normal, block
checksums)

Physical
Position Disk Pool Type RPM Usable
Size Status Size
-----
-----
dparity 1.31.7 1 BSAS 7200 827.7GB
828.0GB (normal)
parity 1.31.6 1 BSAS 7200 827.7GB
828.0GB (normal)
data 1.31.3 1 BSAS 7200 827.7GB
828.0GB (normal)
data 1.31.4 1 BSAS 7200 827.7GB
828.0GB (normal)
```

```

    data      1.31.5                1   BSAS      7200   827.7GB
828.0GB (normal)

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block
checksums)
Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active,
pool0)
RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block
checksums)

```

Physical	Position	Disk	Pool	Type	RPM	Size	Usable
Size	Status						
-----	-----	-----	-----	-----	-----	-----	-----
	dparity	2.30.12	0	BSAS	7200	827.7GB	
828.0GB (normal)	parity	2.30.22	0	BSAS	7200	827.7GB	
828.0GB (normal)	data	2.30.21	0	BSAS	7200	827.7GB	
828.0GB (normal)	data	2.30.20	0	BSAS	7200	827.7GB	
828.0GB (normal)	data	2.30.14	0	BSAS	7200	827.7GB	
828.0GB (normal)							

15 entries were displayed.

5. Physically power off the shelf that you selected.

6. Check the aggregate status again:

```
storage aggregate show
```

```
storage aggregate show -r -node node_A_2 !*root
```

The aggregate with drives on the powered-off shelf should have a "degraded" RAID status, and drives on the affected plex should have a "failed" status, as shown in the following example:

```

cluster_A::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
node_A_1data01_mirrored
           4.15TB    3.40TB   18% online    3 node_A_1
raid_dp,

```

```

mirrored,

normal
node_A_1root
      707.7GB   34.29GB   95% online      1 node_A_1
raid_dp,

mirrored,

normal
node_A_2_data01_mirrored
      4.15TB    4.12TB    1% online      2 node_A_2
raid_dp,

mirror

degraded
node_A_2_data02_unmirrored
      2.18TB    2.18TB    0% online      1 node_A_2
raid_dp,

normal
node_A_2_root
      707.7GB   34.27GB   95% online      1 node_A_2
raid_dp,

mirror

degraded
cluster_A::> storage aggregate show -r -node node_A_2 !*root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirror degraded)
(block checksums)
Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block
checksums)

                                         Usable
Physical
      Position Disk                               Pool Type      RPM      Size
Size Status
-----
      dparity  2.30.3                               0   BSAS       7200  827.7GB
828.0GB (normal)
      parity   2.30.4                               0   BSAS       7200  827.7GB
828.0GB (normal)

```

```

    data      2.30.6          0   BSAS    7200  827.7GB
828.0GB (normal)
    data      2.30.8          0   BSAS    7200  827.7GB
828.0GB (normal)
    data      2.30.5          0   BSAS    7200  827.7GB
828.0GB (normal)

```

Plex: /node_A_2_data01_mirrored/plex4 (offline, failed, inactive, pool1)

RAID Group /node_A_2_data01_mirrored/plex4/rg0 (partial, none checksums)

					Usable
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	FAILED	-	-	-	827.7GB
- (failed)					
parity	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block checksums)

Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active, pool0)

RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block checksums)

					Usable
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	2.30.12	0	BSAS	7200	827.7GB
828.0GB (normal)					
parity	2.30.22	0	BSAS	7200	827.7GB
828.0GB (normal)					
data	2.30.21	0	BSAS	7200	827.7GB
828.0GB (normal)					

```
data      2.30.20      0   BSAS   7200  827.7GB
828.0GB (normal)
data      2.30.14      0   BSAS   7200  827.7GB
828.0GB (normal)
15 entries were displayed.
```

7. Verify that the data is being served and that all volumes are still online:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0,!MDV*
```

```

cluster_A::> vserver show -type data

cluster_A::> vserver show -type data
Admin      Operational Root
Vserver    Type      Subtype   State     State     Volume
Aggregate
-----
-----
SVM1       data      sync-source      running   SVM1_root
node_A_1_data01_mirrored
SVM2       data      sync-source      running   SVM2_root
node_A_1_data01_mirrored

cluster_A::> network interface show -fields is-home false
There are no entries matching your query.

cluster_A::> volume show !vol0,!MDV*
Vserver    Volume      Aggregate    State     Type     Size
Available Used%
-----
-----
SVM1
          SVM1_root
                node_A_1data01_mirrored
                        online    RW      10GB
9.50GB    5%
SVM1
          SVM1_data_vol
                node_A_1data01_mirrored
                        online    RW      10GB
9.49GB    5%
SVM2
          SVM2_root
                node_A_1data01_mirrored
                        online    RW      10GB
9.49GB    5%
SVM2
          SVM2_data_vol
                node_A_2_data02_unmirrored
                        online    RW      1GB
972.6MB   5%

```

8. Physically power on the shelf.

Resynchronization starts automatically.

9. Verify that resynchronization has started:

```
storage aggregate show
```

The affected aggregate should have a RAID status of "resyncing", as shown in the following example:

```
cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
node_A_1_data01_mirrored
      4.15TB      3.40TB   18% online    3 node_A_1
raid_dp,
mirrored,
normal
node_A_1_root
      707.7GB    34.29GB   95% online    1 node_A_1
raid_dp,
mirrored,
normal
node_A_2_data01_mirrored
      4.15TB      4.12TB    1% online    2 node_A_2
raid_dp,
resyncing
node_A_2_data02_unmirrored
      2.18TB      2.18TB    0% online    1 node_A_2
raid_dp,
normal
node_A_2_root
      707.7GB    34.27GB   95% online    1 node_A_2
raid_dp,
resyncing
```

10. Monitor the aggregate to confirm that resynchronization is complete:

```
storage aggregate show
```

The affected aggregate should have a RAID status of "normal", as shown in the following example:

```

cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
node_A_1data01_mirrored
          4.15TB      3.40TB   18% online    3 node_A_1
raid_dp,

mirrored,

normal
node_A_1root
          707.7GB    34.29GB   95% online    1 node_A_1
raid_dp,

mirrored,

normal
node_A_2_data01_mirrored
          4.15TB      4.12TB    1% online    2 node_A_2
raid_dp,

normal
node_A_2_data02_unmirrored
          2.18TB      2.18TB    0% online    1 node_A_2
raid_dp,

normal
node_A_2_root
          707.7GB    34.27GB   95% online    1 node_A_2
raid_dp,

resyncing

```

Remove MetroCluster configurations

If you need to remove the MetroCluster configuration, contact technical support.

Contact NetApp technical support and reference the appropriate guide for your configuration from [How to remove nodes from a MetroCluster configuration - Resolution Guide](#).



You cannot reverse the MetroCluster unconfiguration. This process should only be done with the assistance of technical support. After removing the MetroCluster configuration, all disk connectivity and interconnects should be adjusted to be in a supported state.

Requirements and considerations for ONTAP operations with MetroCluster IP configurations

When using ONTAP in a MetroCluster configuration, you should be aware of certain considerations for licensing, peering to clusters outside the MetroCluster configuration, performing volume operations, NVFAIL operations, and other ONTAP operations.

The ONTAP configuration of the two clusters, including networking, should be identical, because the MetroCluster feature relies on the ability of a cluster to seamlessly serve data for its partner in the event of a switchover.

Licensing considerations

- Both sites should be licensed for the same site-licensed features.
- All nodes should be licensed for the same node-locked features.

SnapMirror consideration

- SnapMirror SVM disaster recovery is only supported on MetroCluster configurations running versions of ONTAP 9.5 or later.

MetroCluster operations in ONTAP System Manager

Depending on your ONTAP version, some MetroCluster-specific operations can be performed using ONTAP System Manager.

To learn more, refer to the [Manage MetroCluster sites with System Manager](#) documentation.

FlexCache support in a MetroCluster configuration

Beginning with ONTAP 9.7, FlexCache volumes are supported on MetroCluster configurations. You should be aware of requirements for manual repeer after switchover or switchback operations.

SVM repeer after switchover when FlexCache origin and cache are within the same MetroCluster site

After a negotiated or unplanned switchover, any SVM FlexCache peering relationship within the cluster must be manually configured.

For example, SVMs vs1 (cache) and vs2 (origin) are on site_A. These SVMs are peered.

After switchover, SVMs vs1-mc and vs2-mc are activated at the partner site (site_B). They must be manually repeer for FlexCache to work using the `vserver peer repeer` command.

SVM repeer after switchover or switchback when a FlexCache destination is on a third cluster and in disconnected mode

For FlexCache relationships to a cluster outside of the MetroCluster configuration, the peering must always be manually reconfigured after a switchover if the involved clusters are in disconnected mode during switchover.

For example:

- One end of the FlexCache (cache_1 on vs1) resides on MetroCluster site_A has one end of the FlexCache
- The other end of the FlexCache (origin_1 on vs2) resides on site_C (not in the MetroCluster configuration)

When switchover is triggered, and if site_A and site_C are not connected, you must manually repeer the SVMs on site_B (the switchover cluster) and site_C using the vserver peer repeer command after the switchover.

When switchback is performed, you must again repeer the SVMs on site_A (the original cluster) and site_C.

Related information

[FlexCache volumes management with the CLI](#)

FabricPool support in MetroCluster configurations

Beginning with ONTAP 9.7, MetroCluster configurations support FabricPool storage tiers.

For general information on using FabricPools, see [Disk and tier \(aggregate\) management](#).

Considerations when using FabricPools

- The clusters must have FabricPool licenses with matching capacity limits.
- The clusters must have IPspaces with matching names.

This can be the default IPspace, or an IP space an administrator has created. This IPspace will be used for FabricPool object store configuration setups.

- For the selected IPspace, each cluster must have an intercluster LIF defined that can reach the external object store.
- SVM migration isn't supported with FabricPool when the source or destination is a MetroCluster cluster.

[Learn more about SVM data mobility.](#)

Configuring an aggregate for use in a mirrored FabricPool



Before you configure the aggregate you must set up object stores as described in "Setting up object stores for FabricPool in a MetroCluster configuration" in [Disk and aggregate management](#).

Steps

To configure an aggregate for use in a FabricPool:

1. Create the aggregate or select an existing aggregate.
2. Mirror the aggregate as a typical mirrored aggregate within the MetroCluster configuration.
3. Create the FabricPool mirror with the aggregate, as described in [Disk and aggregate management](#)

- a. Attach a primary object store.

This object store is physically closer to the cluster.

- b. Add a mirror object store.

This object store is physically further distant to the cluster than the primary object store.

FlexGroup support in MetroCluster configurations

Beginning with ONTAP 9.6 MetroCluster configurations support FlexGroup volumes.

Job schedules in a MetroCluster configuration

In ONTAP 9.3 and later, user-created job schedules are automatically replicated between clusters in a MetroCluster configuration. If you create, modify, or delete a job schedule on a cluster, the same schedule is automatically created on the partner cluster, using Configuration Replication Service (CRS).



System-created schedules are not replicated and you must manually perform the same operation on the partner cluster so that job schedules on both clusters are identical.

Cluster peering from the MetroCluster site to a third cluster

Because the peering configuration is not replicated, if you peer one of the clusters in the MetroCluster configuration to a third cluster outside of that configuration, you must also configure the peering on the partner MetroCluster cluster. This is so that peering can be maintained if a switchover occurs.

The non-MetroCluster cluster must be running ONTAP 8.3 or later. If not, peering is lost if a switchover occurs even if the peering has been configured on both MetroCluster partners.

LDAP client configuration replication in a MetroCluster configuration

An LDAP client configuration created on a storage virtual machine (SVM) on a local cluster is replicated to its partner data SVM on the remote cluster. For example, if the LDAP client configuration is created on the admin SVM on the local cluster, then it is replicated to all the admin data SVMs on the remote cluster. This MetroCluster feature is intentional so that the LDAP client configuration is active on all the partner SVMs on the remote cluster.

Networking and LIF creation guidelines for MetroCluster configurations

You should be aware of how LIFs are created and replicated in a MetroCluster configuration. You must also know about the requirement for consistency so that you can make proper decisions when configuring your network.

Related information

[Network and LIF management](#)

[IPspace object replication and subnet configuration requirements](#)

[Requirements for LIF creation in a MetroCluster configuration](#)

[LIF replication and placement requirements and issues](#)

IPspace object replication and subnet configuration requirements

You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace replication

You must consider the following guidelines while replicating IPspace objects to the partner cluster:

- The IPspace names of the two sites must match.
- IPspace objects must be manually replicated to the partner cluster.

Any storage virtual machines (SVMs) that are created and assigned to an IPspace before the IPspace is replicated will not be replicated to the partner cluster.

Subnet configuration

You must consider the following guidelines while configuring subnets in a MetroCluster configuration:

- Both clusters of the MetroCluster configuration must have a subnet in the same IPspace with the same subnet name, subnet, broadcast domain, and gateway.
- The IP ranges of the two clusters must be different.

In the following example, the IP ranges are different:

```
cluster_A::> network subnet show

IPspace: Default
Subnet
Name      Subnet          Broadcast
-----  -
Domain    Gateway         Total    Ranges
-----  -
subnet1   192.168.2.0/24  Default  192.168.2.1    10/10
192.168.2.11-192.168.2.20

cluster_B::> network subnet show
IPspace: Default
Subnet
Name      Subnet          Broadcast
-----  -
Domain    Gateway         Total    Ranges
-----  -
subnet1   192.168.2.0/24  Default  192.168.2.1    10/10
192.168.2.21-192.168.2.30
```

IPv6 configuration

If IPv6 is configured on one site, IPv6 must be configured on the other site as well.

Related information

[Requirements for LIF creation in a MetroCluster configuration](#)

[LIF replication and placement requirements and issues](#)

Requirements for LIF creation in a MetroCluster configuration

You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

You must consider the following guidelines when creating LIFs:

- Fibre Channel: You must use stretched VSAN or stretched fabrics
- IP/iSCSI: You must use layer 2 stretched network
- ARP broadcasts: You must enable ARP broadcasts between the two clusters
- Duplicate LIFs: You must not create multiple LIFs with the same IP address (duplicate LIFs) in an IPspace
- NFS and SAN configurations: You must use different storage virtual machines (SVMs) for both the unmirrored and mirrored aggregates
- You should create a subnet object before you create a LIF. A subnet object enables ONTAP to determine failover targets on the destination cluster because it has an associated broadcast domain.

Verify LIF creation

You can confirm the successful creation of a LIF in a MetroCluster configuration by running the `metrocluster check lif show` command. If you encounter any issues while creating the LIF, you can use the `metrocluster check lif repair-placement` command to fix the issues.

Related information

[IPspace object replication and subnet configuration requirements](#)

[LIF replication and placement requirements and issues](#)

LIF replication and placement requirements and issues

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

Replication of LIFs to the partner cluster

When you create a LIF on a cluster in a MetroCluster configuration, the LIF is replicated on the partner cluster. LIFs are not placed on a one-to-one name basis. For availability of LIFs after a switchover operation, the LIF placement process verifies that the ports are able to host the LIF based on reachability and port attribute checks.

The system must meet the following conditions to place the replicated LIFs on the partner cluster:

Condition	LIF type: FC	LIF type: IP/iSCSI
-----------	--------------	--------------------

Node identification	ONTAP attempts to place the replicated LIF on the disaster recovery (DR) partner of the node on which it was created. If the DR partner is unavailable, the DR auxiliary partner is used for placement.	ONTAP attempts to place the replicated LIF on the DR partner of the node on which it was created. If the DR partner is unavailable, the DR auxiliary partner is used for placement.
Port identification	ONTAP identifies the connected FC target ports on the DR cluster.	<p>The ports on the DR cluster that are in the same IPspace as the source LIF are selected for a reachability check. If there are no ports in the DR cluster in the same IPspace, the LIF cannot be placed.</p> <p>All of the ports in the DR cluster that are already hosting a LIF in the same IPspace and subnet are automatically marked as reachable; and can be used for placement. These ports are not included in the reachability check.</p>
Reachability check	Reachability is determined by checking for the connectivity of the source fabric WWN on the ports in the DR cluster. If the same fabric is not present at the DR site, the LIF is placed on a random port on the DR partner.	<p>Reachability is determined by the response to an Address Resolution Protocol (ARP) broadcast from each previously identified port on the DR cluster to the source IP address of the LIF to be placed. For reachability checks to succeed, ARP broadcasts must be allowed between the two clusters.</p> <p>Each port that receives a response from the source LIF will be marked as possible for placement.</p>
Port selection	ONTAP categorizes the ports based on attributes such as adapter type and speed, and then selects the ports with matching attributes. If no ports with matching attributes are found, the LIF is placed on a random connected port on the DR partner.	<p>From the ports that are marked as reachable during the reachability check, ONTAP prefers ports that are in the broadcast domain that is associated with the subnet of the LIF. If there are no network ports available on the DR cluster that are in the broadcast domain that is associated with the subnet of the LIF, then ONTAP selects ports that have reachability to the source LIF.</p> <p>If there are no ports with reachability to the source LIF, a port is selected from the broadcast domain that is associated with the subnet of the source LIF, and if no such broadcast domain exists, a random port is selected.</p> <p>ONTAP categorizes the ports based on attributes such as adapter type, interface type, and speed, and then selects the ports with matching attributes.</p>
LIF placement	From the reachable ports, ONTAP selects the least loaded port for placement.	From the selected ports, ONTAP selects the least loaded port for placement.

Placement of replicated LIFs when the DR partner node is down

When an iSCSI or FC LIF is created on a node whose DR partner has been taken over, the replicated LIF is placed on the DR auxiliary partner node. After a subsequent giveback operation, the LIFs are not automatically moved to the DR partner. This can lead to LIFs being concentrated on a single node in the partner cluster. During a MetroCluster switchover operation, subsequent attempts to map LUNs belonging to the storage virtual machine (SVM) fail.

You should run the `metrocluster check lif show` command after a takeover operation or giveback operation to verify that the LIF placement is correct. If errors exist, you can run the `metrocluster check lif repair-placement` command to resolve the issues.

LIF placement errors

LIF placement errors that are displayed by the `metrocluster check lif show` command are retained after a switchover operation. If the `network interface modify`, `network interface rename`, or `network interface delete` command is issued for a LIF with a placement error, the error is removed and does not appear in the output of the `metrocluster check lif show` command.

LIF replication failure

You can also check whether LIF replication was successful by using the `metrocluster check lif show` command. An EMS message is displayed if LIF replication fails.

You can correct a replication failure by running the `metrocluster check lif repair-placement` command for any LIF that fails to find a correct port. You should resolve any LIF replication failures as soon as possible to verify the availability of LIF during a MetroCluster switchover operation.



Even if the source SVM is down, LIF placement might proceed normally if there is a LIF belonging to a different SVM in a port with the same IPspace and network in the destination SVM.

Related information

[IPspace object replication and subnet configuration requirements](#)

[Requirements for LIF creation in a MetroCluster configuration](#)

Volume creation on a root aggregate

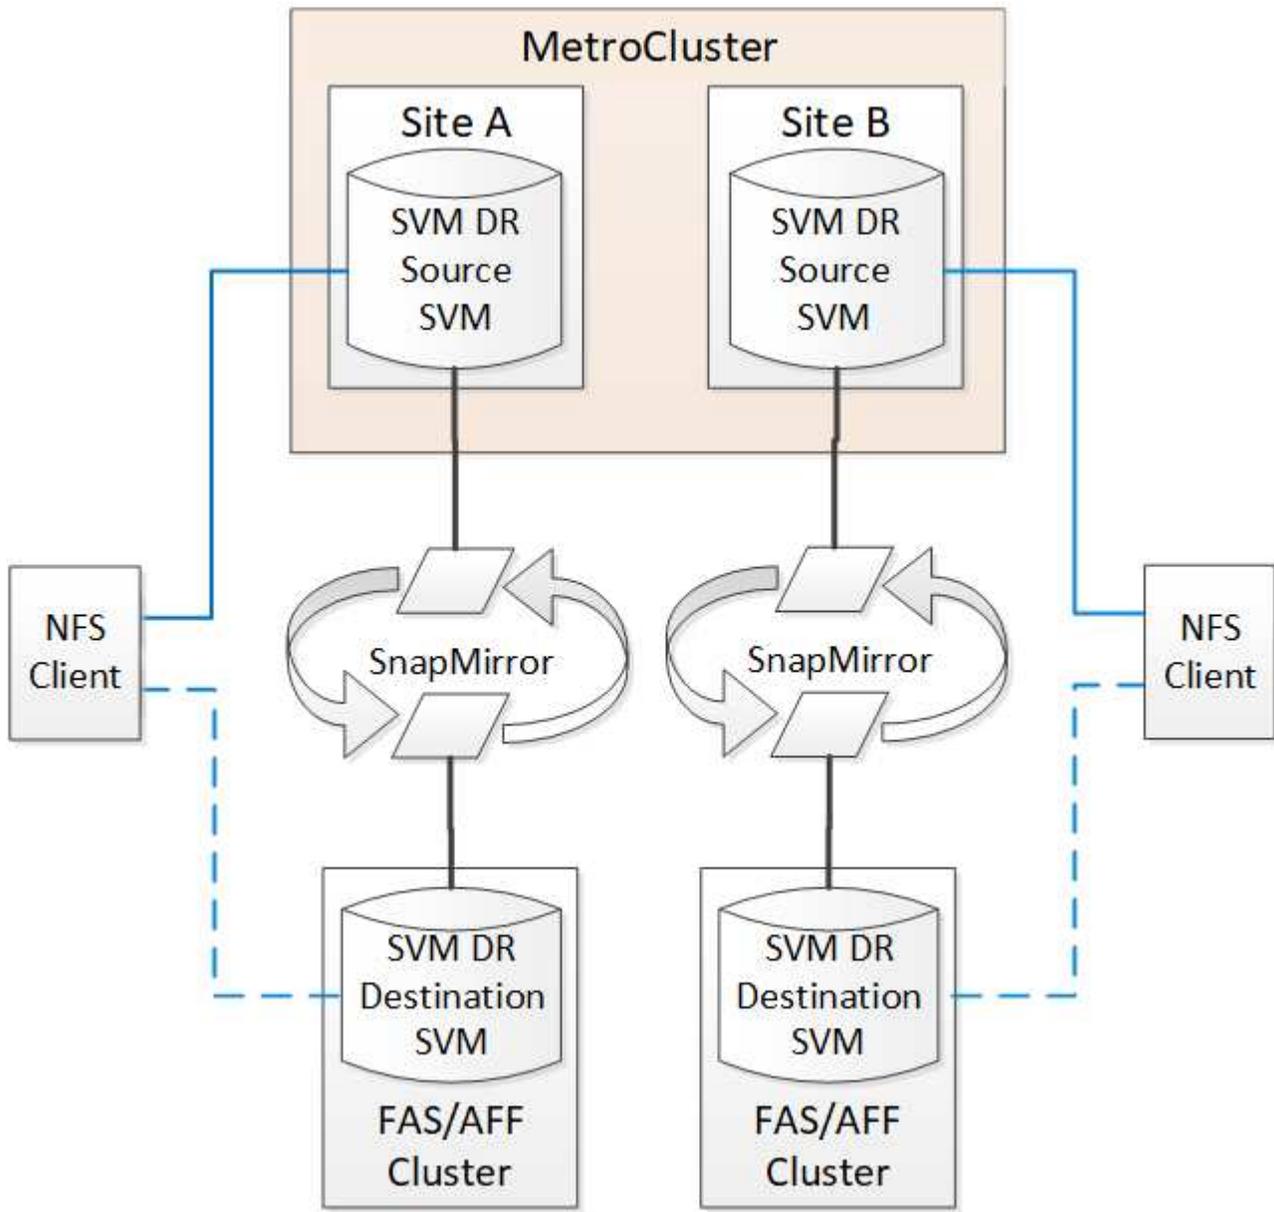
The system does not allow the creation of new volumes on the root aggregate (an aggregate with an HA policy of CFO) of a node in a MetroCluster configuration.

Because of this restriction, root aggregates cannot be added to an SVM using the `vserver add-aggregates` command.

SVM disaster recovery in a MetroCluster configuration

Beginning with ONTAP 9.5, active storage virtual machines (SVMs) in a MetroCluster configuration can be used as sources with the SnapMirror SVM disaster recovery feature. The destination SVM must be on the third cluster outside of the MetroCluster configuration.

Beginning with ONTAP 9.11.1, both sites within a MetroCluster configuration can be the source for an SVM DR relationship with a FAS or AFF destination cluster as shown in the following image.



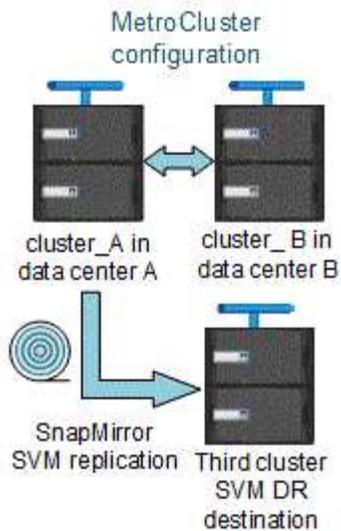
You should be aware of the following requirements and limitations of using SVMs with SnapMirror disaster recovery:

- Only an active SVM within a MetroCluster configuration can be the source of an SVM disaster recovery relationship.

A source can be a sync-source SVM before switchover or a sync-destination SVM after switchover.

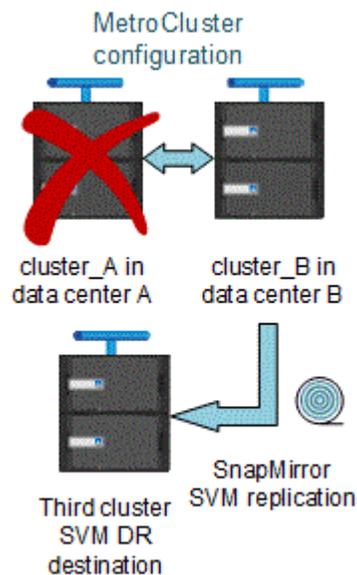
- When a MetroCluster configuration is in a steady state, the MetroCluster sync-destination SVM cannot be the source of an SVM disaster recovery relationship, since the volumes are not online.

The following image shows the SVM disaster recovery behavior in a steady state:



- When the sync-source SVM is the source of an SVM DR relationship, the source SVM DR relationship information is replicated to the MetroCluster partner.

This enables the SVM DR updates to continue after a switchover as shown in the following image:



- During the switchover and switchback processes, replication to the SVM DR destination might fail.

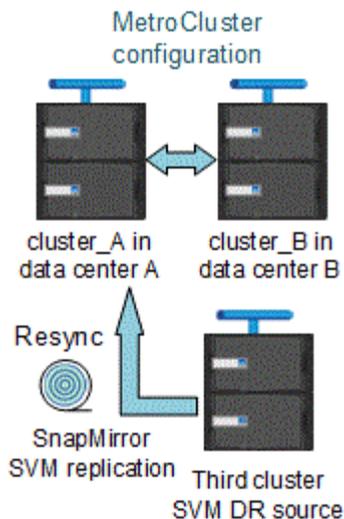
However, after the switchover or switchback process completes, the next SVM DR scheduled updates will succeed.

See “Replicating the SVM configuration” in [Data protection](#) for details on configuring an SVM DR relationship.

SVM resynchronization at a disaster recovery site

During resynchronization, the storage virtual machines (SVMs) disaster recovery (DR) source on the MetroCluster configuration is restored from the destination SVM on the non-MetroCluster site.

During resynchronization, the source SVM (cluster_A) temporarily acts as a destination SVM as shown in the following image:



If an unplanned switchover occurs during resynchronization

Unplanned switchovers that occur during the resynchronization will halt the resynchronization transfer. If an unplanned switchover occurs, the following conditions are true:

- The destination SVM on the MetroCluster site (which was a source SVM prior to resynchronization) remains as a destination SVM. The SVM at the partner cluster will continue to retain its subtype and remain inactive.
- The SnapMirror relationship must be re-created manually with the sync-destination SVM as the destination.
- The SnapMirror relationship does not appear in the SnapMirror show output after a switchover at the survivor site unless a SnapMirror create operation is executed.

Performing switchback after an unplanned switchover during resynchronization

To successfully perform the switchback process, the resynchronization relationship must be broken and deleted. Switchback is not permitted if there are any SnapMirror DR destination SVMs in the MetroCluster configuration or if the cluster has an SVM of subtype "dp-destination".

Output for the storage aggregate plex show command is indeterminate after a MetroCluster switchover

When you run the storage aggregate plex show command after a MetroCluster switchover, the status of plex0 of the switched over root aggregate is indeterminate and is displayed as failed. During this time, the switched over root is not updated. The actual status of this plex can only be determined after the MetroCluster healing phase.

Modifying volumes to set the NVFAIL flag in case of switchover

You can modify a volume so that the NVFAIL flag is set on the volume in the event of a MetroCluster switchover. The NVFAIL flag causes the volume to be fenced off from any modification. This is required for volumes that need to be handled as if committed writes to the volume were lost after the switchover.



In ONTAP versions earlier than 9.0, the NVFAIL flag is used for each switchover. In ONTAP 9.0 and later versions, the unplanned switchover (USO) is used.

Step

1. Enable MetroCluster configuration to trigger NVFAIL on switchover by setting the `vol -dr-force -nvfail` parameter to on:

```
vol modify -vserver vserver-name -volume volume-name -dr-force-nvfail on
```

How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring in a MetroCluster IP configuration

The Active IQ Unified Manager and ONTAP System Manager can be used for GUI management of the clusters and monitoring the configuration.

Each node has ONTAP System Manager pre-installed. To load System Manager, enter the cluster management LIF address as the URL in a web browser that has connectivity to the node.

You can also use Active IQ Unified Manager to monitor the MetroCluster configuration.

Related information

[Active IQ Unified Manager Documentation](#)

Synchronize the system time using NTP in a MetroCluster IP configuration

Each cluster needs its own Network Time Protocol (NTP) server to synchronize the time between the nodes and their clients.

About this task

- You cannot modify the time zone settings for a failed node or the partner node after a takeover occurs.
- Each cluster in the MetroCluster IP configuration should have its own separate NTP server or servers used by the nodes and IP switches at that MetroCluster site.
- If you are using the MetroCluster Tiebreaker or ONTAP Mediator, it should also have its own separate NTP server.
- This procedure shows how to configure the NTP after you have already set up the MetroCluster IP clusters. If you used System Manager to configure the clusters, you should already have configured the NTP servers as part of cluster setup. See [Set up a MetroCluster IP site](#) for details.

Depending on your ONTAP version, you can configure the NTP from the **Cluster** or **Insights** tab in the System Manager UI.

Cluster

In System Manager, you can configure the NTP from the **Cluster** tab using two different options, depending on your ONTAP version:

ONTAP 9.8 or later:

Use the following steps to synchronize the NTP from the **Cluster** tab in ONTAP 9.8 or later.

Steps

1. Go to **Cluster > Overview**
2. Then select the  **More** option and select **Edit**.
3. In the **Edit Cluster Details** window, select the **+Add** option below NTP Servers.
4. Add the name, location, and specify the IP address of the time server.
5. Then, select **Save**.
6. Repeat the steps for any additional time servers.

ONTAP 9.11.1 or later:

Use the following steps to synchronize the NTP from the **Insights** window in the **Cluster** tab in ONTAP 9.11.1 or later.

Steps

1. Go to **Cluster > Overview**
2. Scroll down to the **Insights** window on the page, locate **Too few NTP servers are configured**, and then select **Fix It**.
3. Specify the IP address of the time server, and then select **Save**.
4. Repeat the previous step for any additional time servers.

Insights

In ONTAP 9.11.1 or later, you can also configure the NTP by using the **Insights** tab in System Manager:

Steps

1. Go to the **Insights** tab in the System Manager UI.
2. Scroll down to **Too few NTP servers are configured** and select **Fix It**.
3. Specify the IP address of the time server, and then select **Save**.
4. Repeat the previous step for any additional time servers.

Where to find additional information about MetroCluster IP

You can learn more about MetroCluster configuration.

MetroCluster and miscellaneous information

Information	Subject
-------------	---------

MetroCluster IP Solution Architecture and Design, TR-4689	<ul style="list-style-type: none"> • A technical overview of the MetroCluster IP configuration and operation. • Best practices for a MetroCluster IP configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster management	<ul style="list-style-type: none"> • Understanding the MetroCluster configuration • Switchover, healing, and switchback
Disaster Recovery	<ul style="list-style-type: none"> • Disaster recovery • Forced switchover • Recovery from a multi-controller or storage failure
MetroCluster Maintenance	<ul style="list-style-type: none"> • Guidelines for maintenance in a MetroCluster FC configuration • Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches • Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Replacing hardware at a disaster site in a fabric-attached or stretch MetroCluster FC configuration • Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration. • Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.

MetroCluster Upgrade and Expansion	<ul style="list-style-type: none"> • Upgrading or refreshing a MetroCluster configuration • Expanding a MetroCluster configuration by adding additional nodes
MetroCluster Transition	<ul style="list-style-type: none"> • Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration
MetroCluster Upgrade, Transition, and Expansion	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
ONTAP Hardware Systems Documentation Note: The standard storage shelf maintenance procedures can be used with MetroCluster IP configurations.	<ul style="list-style-type: none"> • Hot-adding a disk shelf • Hot-removing a disk shelf
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems
ONTAP concepts	<ul style="list-style-type: none"> • How mirrored aggregates work

Install a stretch MetroCluster configuration

Overview

To install your stretch MetroCluster configuration, you must perform a number of procedures in the correct order.

- [Prepare for the installation and understand all requirements](#)
- [Choose the correct installation procedure](#)
- Cable the components
 - [Two-node SAS-attached configuration](#)
 - [Two-node bridge-attached configuration](#)
- [Configure the software](#)
- [Test the configuration](#)

Prepare for the MetroCluster installation

Differences between the ONTAP MetroCluster configurations

The various MetroCluster configurations have key differences in the required components.

In all configurations, each of the two MetroCluster sites are configured as an ONTAP cluster. In a two-node MetroCluster configuration, each node is configured as a single-node cluster.

Feature	IP configurations	Fabric attached configurations		Stretch configurations	
		Four- or eight-node	Two-node	Two-node bridge-attached	Two-node direct-attached
Number of controllers	Four or eight ¹	Four or eight	Two	Two	Two
Uses an FC switch storage fabric	No	Yes	Yes	No	No
Uses an IP switch storage fabric	Yes	No	No	No	No
Uses FC-to-SAS bridges	No	Yes	Yes	Yes	No

Uses direct-attached SAS storage	Yes (local attached only)	No	No	No	Yes
Supports ADP	Yes (beginning with ONTAP 9.4)	No	No	No	No
Supports local HA	Yes	Yes	No	No	No
Supports ONTAP automatic unplanned switchover (AUSO)	No	Yes	Yes	Yes	Yes
Supports unmirrored aggregates	Yes (beginning with ONTAP 9.8)	Yes	Yes	Yes	Yes
Supports ONTAP Mediator	Yes (beginning with ONTAP 9.7)	No	No	No	No
Supports MetroCluster Tiebreaker	Yes (not in combination with ONTAP Mediator)	Yes	Yes	Yes	Yes
Supports All SAN Arrays	Yes	Yes	Yes	Yes	Yes

Notes

- Review the following considerations for eight-node MetroCluster IP configurations:
 - Eight-node configurations are supported beginning with ONTAP 9.9.1.
 - Only NetApp-validated MetroCluster switches (ordered from NetApp) are supported.
 - Configurations using IP-routed (layer 3) backend connections are not supported.

Support for All SAN Array systems in MetroCluster configurations

Some of the All SAN Arrays (ASAs) are supported in MetroCluster configurations. In the MetroCluster documentation, the information for AFF models applies to the corresponding ASA system. For example, all cabling and other information for the AFF A400 system also applies to the ASA AFF A400 system.

Supported platform configurations are listed in the [NetApp Hardware Universe](#).

Cluster peering

Each MetroCluster site is configured as a peer to its partner site. You must be familiar with the prerequisites and guidelines for configuring the peering relationships. This is important when deciding on whether to use shared or dedicated ports for those relationships.

Related information

[Cluster and SVM peering express configuration](#)

Prerequisites for cluster peering

Before you set up cluster peering, you should confirm that connectivity between port, IP address, subnet, firewall, and cluster-naming requirements are met.

Connectivity requirements

Every intercluster LIF on the local cluster must be able to communicate with every intercluster LIF on the remote cluster.

Although it is not required, it is typically simpler to configure the IP addresses used for intercluster LIFs in the same subnet. The IP addresses can reside in the same subnet as data LIFs, or in a different subnet. The subnet used in each cluster must meet the following requirements:

- The subnet must have enough IP addresses available to allocate to one intercluster LIF per node.

For example, in a four-node cluster, the subnet used for intercluster communication must have four available IP addresses.

Each node must have an intercluster LIF with an IP address on the intercluster network.

Intercluster LIFs can have an IPv4 address or an IPv6 address.



ONTAP 9 enables you to migrate your peering networks from IPv4 to IPv6 by optionally allowing both protocols to be present simultaneously on the intercluster LIFs. In earlier releases, all intercluster relationships for an entire cluster were either IPv4 or IPv6. This meant that changing protocols was a potentially disruptive event.

Port requirements

You can use dedicated ports for intercluster communication, or share ports used by the data network. Ports must meet the following requirements:

- All ports used to communicate with a given remote cluster must be in the same IPspace.

You can use multiple IPspaces to peer with multiple clusters. Pair-wise full-mesh connectivity is required only within an IPspace.

- The broadcast domain used for intercluster communication must include at least two ports per node so that intercluster communication can fail over from one port to another port.

Ports added to a broadcast domain can be physical network ports, VLANs, or interface groups (ifgrps).

- All ports must be cabled.
- All ports must be in a healthy state.
- The MTU settings of the ports must be consistent.

Firewall requirements

Firewalls and the intercluster firewall policy must allow the following protocols:

- ICMP service
- TCP to the IP addresses of all the intercluster LIFs over the ports 10000, 11104, and 11105
- Bidirectional HTTPS between the intercluster LIFs

The default intercluster firewall policy allows access through the HTTPS protocol and from all IP addresses (0.0.0.0/0). You can modify or replace the policy if necessary.

Considerations when using dedicated ports

When determining whether using a dedicated port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether using a dedicated port is the best intercluster network solution:

- If the amount of available WAN bandwidth is similar to that of the LAN ports, and the replication interval is such that replication occurs while regular client activity exists, then you should dedicate Ethernet ports for intercluster replication to avoid contention between replication and the data protocols.
- If the network utilization generated by the data protocols (CIFS, NFS, and iSCSI) is such that the network utilization is above 50 percent, then dedicate ports for replication to allow for nondegraded performance if a node failover occurs.
- When physical 10 GbE or faster ports are used for data and replication, you can create VLAN ports for replication and dedicate the logical ports for intercluster replication.

The bandwidth of the port is shared between all VLANs and the base port.

- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.

Considerations when sharing data ports

When determining whether sharing a data port for intercluster replication is the correct intercluster network solution, you should consider configurations and requirements such as LAN type, available WAN bandwidth, replication interval, change rate, and number of ports.

Consider the following aspects of your network to determine whether sharing data ports is the best intercluster connectivity solution:

- For a high-speed network, such as a 40-Gigabit Ethernet (40-GbE) network, a sufficient amount of local LAN bandwidth might be available to perform replication on the same 40-GbE ports that are used for data access.

In many cases, the available WAN bandwidth is far less than the 10 GbE LAN bandwidth.

- All nodes in the cluster might have to replicate data and share the available WAN bandwidth, making data port sharing more acceptable.
- Sharing ports for data and replication eliminates the extra port counts required to dedicate ports for replication.
- The maximum transmission unit (MTU) size of the replication network will be the same size as that used on the data network.
- Consider the data change rate and replication interval and whether the amount of data, that must be replicated on each interval, requires enough bandwidth. This might cause contention with data protocols if sharing data ports.
- When data ports for intercluster replication are shared, the intercluster LIFs can be migrated to any other intercluster-capable port on the same node to control the specific data port that is used for replication.

Considerations when using unmirrored aggregates

Considerations when using unmirrored aggregates

If your configuration includes unmirrored aggregates, you must be aware of potential access issues that follow switchover operations.

Considerations for unmirrored aggregates when doing maintenance requiring power shutdown

If you are performing a negotiated switchover for maintenance reasons requiring site-wide power shutdown, you should first manually take offline any unmirrored aggregates owned by the disaster site.

If you do not take any unmirrored aggregates offline, nodes at the surviving site might go down due to multi-disk panics. This could occur if switched over unmirrored aggregates go offline, or are missing, because of the loss of connectivity to storage at the disaster site. This is the result of a power shutdown or a loss of ISLs.

Considerations for unmirrored aggregates and hierarchical namespaces

If you are using hierarchical namespaces, you should configure the junction path so that all of the volumes in that path are either on mirrored aggregates only or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates in the junction path might prevent access to the unmirrored aggregates after the switchover operation.

Considerations for unmirrored aggregates and CRS metadata volume and data SVM root volumes

The configuration replication service (CRS) metadata volume and data SVM root volumes must be on a mirrored aggregate. You cannot move these volumes to an unmirrored aggregate. If they are on an unmirrored aggregate, negotiated switchover and switchback operations are vetoed. The MetroCluster check command provides a warning if this is the case.

Considerations for unmirrored aggregates and SVMs

SVMs should be configured on mirrored aggregates only, or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates can result in a switchover operation that exceeds 120 seconds and result in a data outage if the unmirrored aggregates do not come online.

Considerations for unmirrored aggregates and SAN

In ONTAP versions prior to 9.9.1, a LUN should not be located on an unmirrored aggregate. Configuring a LUN on an unmirrored aggregate can result in a switchover operation that exceeds 120 seconds and a data outage.

Firewall usage at MetroCluster sites

Considerations for firewall usage at MetroCluster sites

If you are using a firewall at a MetroCluster site, you must ensure access for required ports.

The following table shows TCP/UDP port usage in an external firewall positioned between two MetroCluster sites.

Traffic type	Port/services
Cluster peering	11104 / TCP 11105 / TCP
ONTAP System Manager	443 / TCP
MetroCluster IP intercluster LIFs	65200 / TCP 10006 / TCP and UDP
Hardware assist	4444 / TCP

Choosing the correct installation procedure for your configuration

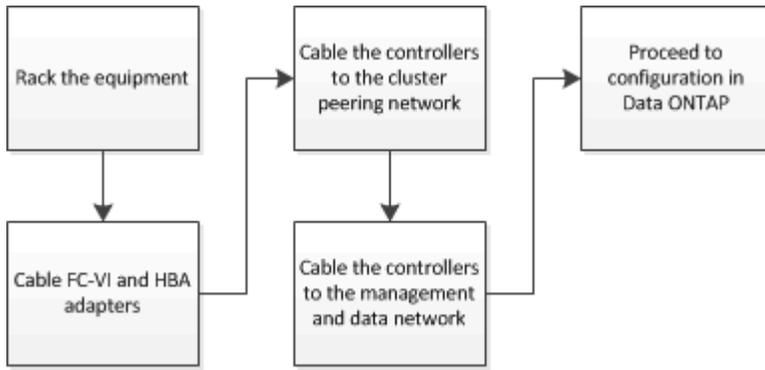
You must choose the correct installation procedure based on how the storage controllers connect to storage shelves.

For this installation type...	Use these procedures...
Two-node stretch configuration with FC-to SAS bridges	<ol style="list-style-type: none">Cabling a two-node bridge-attached stretch MetroCluster configurationConfiguring the MetroCluster software in ONTAP
Two-node stretch configuration with direct-attached SAS cabling	<ol style="list-style-type: none">Cabling a two-node SAS-attached stretch MetroCluster configurationConfiguring the MetroCluster software in ONTAP

Cable a two-node SAS-attached stretch MetroCluster configuration

Cabling a two-node SAS-attached stretch MetroCluster configuration

The MetroCluster components must be physically installed, cabled, and configured at both geographic sites.



Parts of a two-node SAS-attached stretch MetroCluster configuration

The two-node MetroCluster SAS-attached configuration requires a number of parts, including two single-node clusters in which the storage controllers are directly connected to the storage using SAS cables.

The MetroCluster configuration includes the following key hardware elements:

- Storage controllers

The storage controllers connect directly to the storage using SAS cables.

Each storage controller is configured as a DR partner to a storage controller on the partner site.

- Copper SAS cables can be used for shorter distances.
- Optical SAS cables can be used for longer distances.

[NetApp Interoperability Matrix Tool](#)

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- Cluster peering network

The cluster peering network provides connectivity for mirroring of the storage virtual machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

Required MetroCluster hardware components and naming guidelines for two-node SAS-attached stretch configurations

The MetroCluster configuration requires a variety of hardware components. For convenience and clarity, standard names for components are used throughout the MetroCluster documentation. One site is referred to as Site A and the other site is referred to as Site B.

Supported software and hardware

The hardware and software must be supported for the MetroCluster FC configuration.

[NetApp Hardware Universe](#)

When using AFF systems, all controller modules in the MetroCluster configuration must be configured as AFF systems.

Hardware redundancy in the MetroCluster configuration

Because of the hardware redundancy in the MetroCluster configuration, there are two of each components at each site. The sites are arbitrarily assigned the letters A and B and the individual components are arbitrarily assigned the numbers 1 and 2.

Two single-node ONTAP clusters

The SAS-attached stretch MetroCluster configuration requires two single-node ONTAP clusters.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A
- Site B: cluster_B

Two storage controller modules

The SAS-attached stretch MetroCluster configuration requires two storage controller modules.

- Naming must be unique within the MetroCluster configuration.
- All controller modules in the MetroCluster configuration must be running the same version of ONTAP.
- All controller modules in a DR group must be of the same model.
- All controller modules in a DR group must use the same FC-VI configuration.

Some controller modules support two options for FC-VI connectivity:

- Onboard FC-VI ports
- An FC-VI card in slot 1

A mix of one controller module using onboard FC-VI ports and another using an add-on FC-VI card is not supported. For example, if one node uses onboard FC-VI configuration, then all other nodes in the DR group must use onboard FC-VI configuration as well.

Example names:

- Site A: controller_A_1
- Site B: controller_B_1

At least four SAS disk shelves (recommended)

The SAS-attached stretch MetroCluster configuration requires at least two SAS disk shelves. Four SAS disk shelves is recommended.

Two shelves are recommended at each site to allow disk ownership on a per-shelf basis. A minimum of one shelf at each site is supported.

Example names:

- Site A:
 - shelf_A_1_1
 - shelf_A_1_2
- Site B:
 - shelf_B_1_1
 - shelf_B_1_2

Install and cable MetroCluster components for two-node SAS-attached stretch configurations

Installing and cabling MetroCluster components for two-node SAS-attached stretch configurations

The storage controllers must be cabled to the storage media and to each other. The storage controllers must also be cabled to the data and management network.

Before you begin any procedure in this document

The following overall requirements must be met before completing this task:

- Prior to installation you must have familiarized yourself with the considerations and best practices for installing and cabling disk shelves for your disk shelf model.
- All MetroCluster components must be supported.

[NetApp Interoperability Matrix Tool](#)

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. Use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

About this task

- The terms node and controller are used interchangeably.

Racking the hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

This task must be performed on both MetroCluster sites.

Steps

1. Plan the positioning of the MetroCluster components.

The amount of rack space needed depends on the platform model of the storage controllers, the switch types, and the number of disk shelf stacks in your configuration.

2. Using standard shop practices for working with electrical equipment make sure you are properly grounded.
3. Install the storage controllers in the rack or cabinet.

[ONTAP Hardware Systems Documentation](#)

4. Install the disk shelves, daisy-chain the disk shelves in each stack, power them on, and set the shelf IDs.

See the appropriate guide for your disk shelf model for information about daisy-chaining disk shelves and setting shelf IDs.



Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites). When manually setting shelf IDs, you must power-cycle the disk shelf.

Cabling the controllers to each other and the storage shelves

The controller FC-VI adapters must be cabled directly to each other. The controller SAS ports must be cabled to both the remote and local storage stacks.

This task must be performed at both MetroCluster sites.

Steps

1. Cable the FC-VI ports.

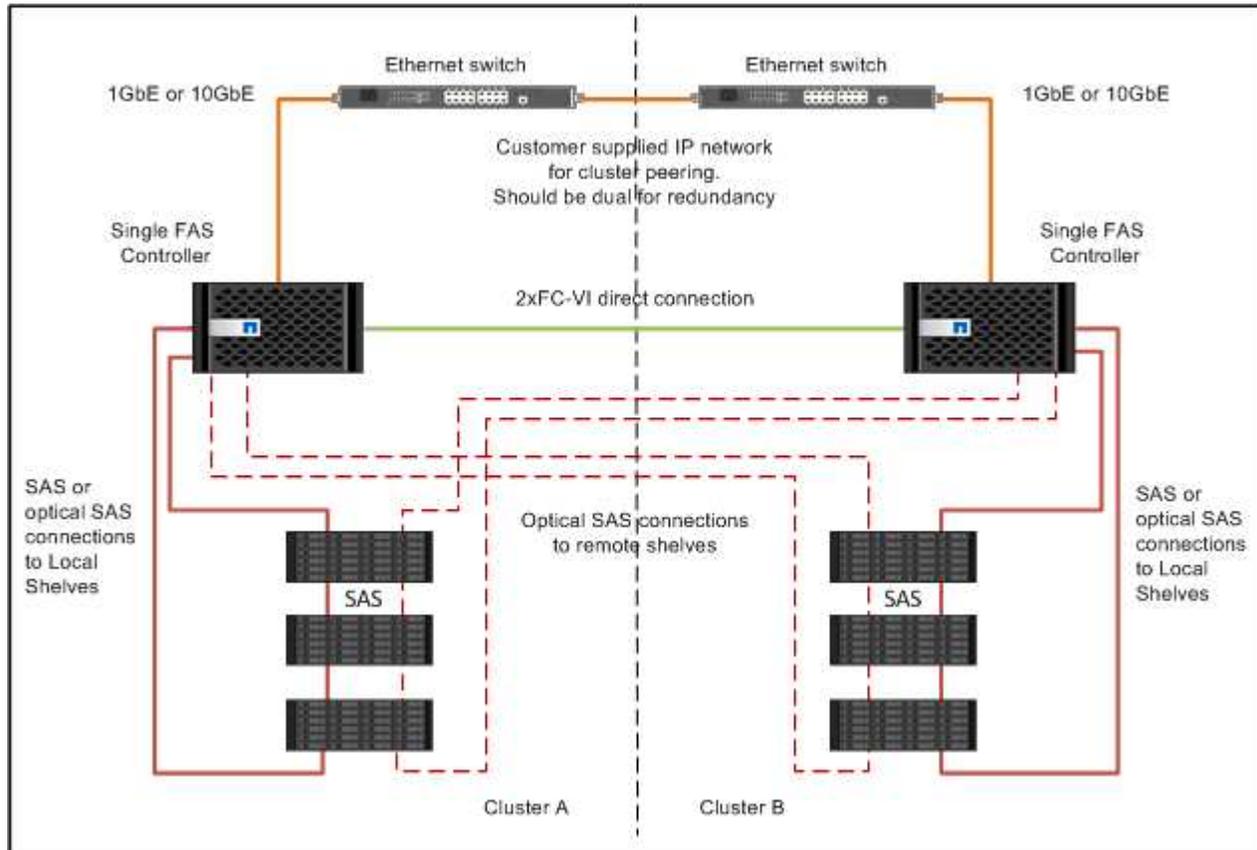


The above illustration is a typical representative cable connection. The specific FC-VI ports will vary by controller module.

- FAS8200 and AFF A300 controller modules can be ordered with one of two options for FC-VI connectivity:
 - Onboard ports 0e and 0f are configured in FC-VI mode.
 - Ports 1a and 1b on an FC-VI card go in slot 1.
- AFF A700 and FAS9000 storage systems controller modules use four FC-VI ports each.
- AFF A400 and FAS8300 storage system controller modules use FC-VI ports 2a and 2b.

2. Cable the SAS ports.

The following illustration shows the connections. Your port usage might be different depending on the available SAS and FC-VI ports on the controller module.



Cabling the cluster peering connections

You must cable the controller module ports used for cluster peering so that they have connectivity with the cluster on their partner site.

This task must be performed on each controller module in the MetroCluster configuration.

At least two ports on each controller module should be used for cluster peering.

The recommended minimum bandwidth for the ports and network connectivity is 1 GbE.

Steps

1. Identify and cable at least two ports for cluster peering and verify they have network connectivity with the partner cluster.

Cluster peering can be done on dedicated ports or on data ports. Using dedicated ports provides a higher throughput for the cluster peering traffic.

[Cluster and SVM peering express configuration](#)

Cabling the management and data connections

You must cable the management and data ports on each storage controller to the site networks.

This task must be repeated for each new controller at both MetroCluster sites.

You can connect the controller and cluster switch management ports to existing switches in your network. In addition you can connect controller to new dedicated network switches such as NetApp CN1601 cluster management switches.

Steps

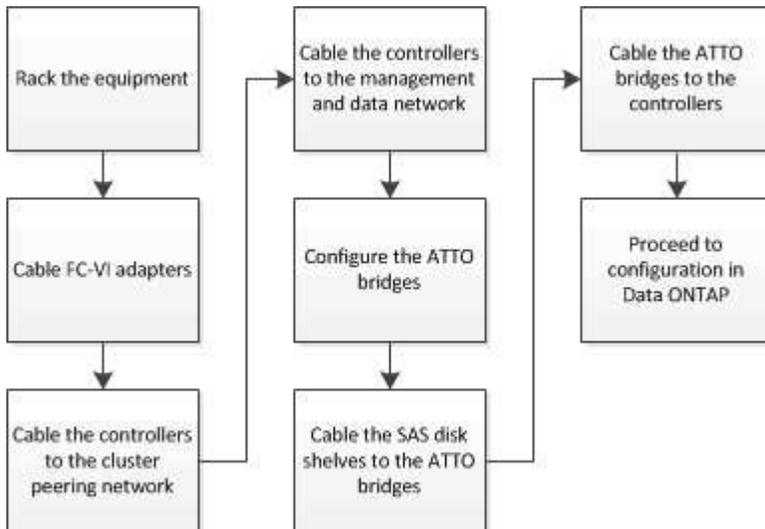
1. Cable the controller's management and data ports to the management and data networks at the local site.

[ONTAP Hardware Systems Documentation](#)

Cable a two-node bridge-attached stretch MetroCluster configuration

Cabling a two-node bridge-attached stretch MetroCluster configuration

The MetroCluster components must be physically installed, cabled, and configured at both geographic sites.



Parts of a two-node bridge-attached stretch MetroCluster configuration

As you plan your MetroCluster configuration, you should understand the parts of the configuration and how they work together.

The MetroCluster configuration includes the following key hardware elements:

- Storage controllers

The storage controllers are not connected directly to the storage but connected to FC-to-SAS bridges. The storage controllers are connected to each other by FC cables between each controller's FC-VI adapters.

Each storage controller is configured as a DR partner to a storage controller on the partner site.

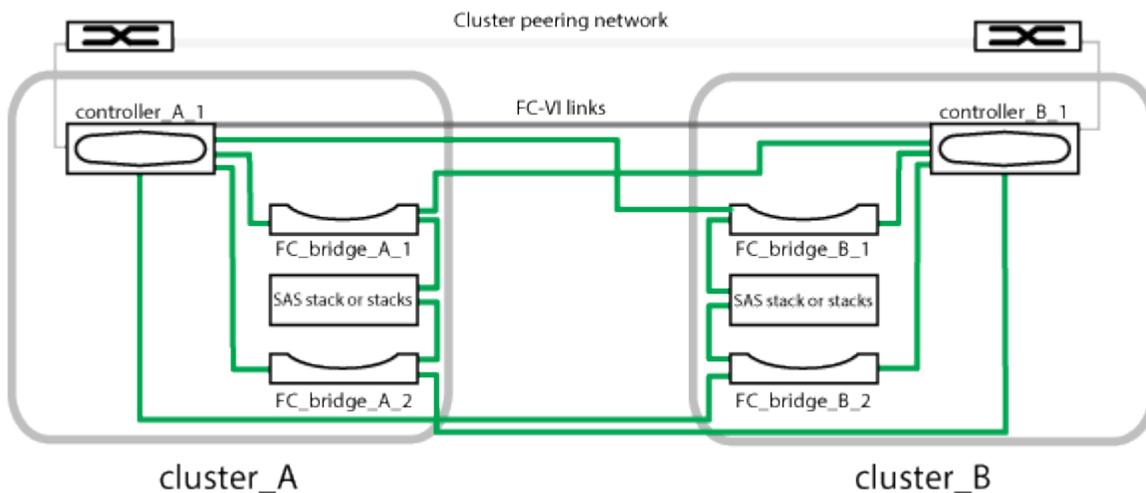
- FC-to-SAS bridges

The FC-to-SAS bridges connect the SAS storage stacks to the FC initiator ports on the controllers, providing bridging between the two protocols.

- Cluster peering network

The cluster peering network provides connectivity for mirroring of the storage virtual machine (SVM) configuration. The configuration of all SVMs on one cluster is mirrored to the partner cluster.

The following illustration shows a simplified view of the MetroCluster configuration. For some connections, a single line represents multiple, redundant connections between the components. Data and management network connections are not shown.



- The configuration consists of two single-node clusters.
- Each site has one or more stacks of SAS storage.



SAS shelves in MetroCluster configurations are not supported with ACP cabling.

Additional storage stacks are supported, but only one is shown at each site.

Required MetroCluster hardware components and naming conventions for two-node bridge-attached stretch configurations

When planning your MetroCluster configuration, you must understand the required and supported hardware and software components. For convenience and clarity, you should also understand the naming conventions used for components in examples throughout the documentation. For example, one site is referred to as Site A and the other site is referred to as Site B.

Supported software and hardware

The hardware and software must be supported for the MetroCluster FC configuration.

When using AFF systems, all controller modules in the MetroCluster configuration must be configured as AFF systems.

Hardware redundancy in the MetroCluster configuration

Because of the hardware redundancy in the MetroCluster configuration, there are two of each component at each site. The sites are arbitrarily assigned the letters A and B and the individual components are arbitrarily assigned the numbers 1 and 2.

Requirement for two single-node ONTAP clusters

The bridge-attached stretch MetroCluster configuration requires two single-node ONTAP clusters.

Naming must be unique within the MetroCluster configuration.

Example names:

- Site A: cluster_A
- Site B: cluster_B

Requirement for two storage controller modules

The bridge-attached stretch MetroCluster configuration requires two storage controller modules.

The controllers must meet the following requirements:

- Naming must be unique within the MetroCluster configuration.
- All controller modules in the MetroCluster configuration must be running the same version of ONTAP.
- All controller modules in a DR group must be of the same model.
- All controller modules in a DR group must use the same FC-VI configuration.

Some controller modules support two options for FC-VI connectivity:

- Onboard FC-VI ports
- An FC-VI card in slot 1

A mix of one controller module using onboard FC-VI ports and another using an add-on FC-VI card is not supported. For example, if one node uses onboard FC-VI configuration, then all other nodes in the DR group must use onboard FC-VI configuration as well.

Example names:

- Site A: controller_A_1
- Site B: controller_B_1

Requirement for FC-to-SAS bridges

The bridge-attached stretch MetroCluster configuration requires two or more FC-to-SAS bridges at each site.

These bridges connect the SAS disk shelves to the controller modules.



FibreBridge 6500N bridges are not supported in configurations running ONTAP 9.8 and later.

- FibreBridge 7600N and 7500N bridges support up to four SAS stacks.
- Each stack can use different models of IOM, but all shelves within a stack must use the same model.

The supported IOM models depend on the ONTAP version you are running.

- Naming must be unique within the MetroCluster configuration.

The suggested names used as examples in this procedure identify the controller module that the bridge connects to and the port.

Example names:

- Site A:
 - *bridge_A_1_port-number*
 - *bridge_A_2_port-number*
- Site B:
 - *bridge_B_1_port-number*
 - *bridge_B_2_port-number*

Requirement for at least four SAS shelves (recommended)

The bridge-attached stretch MetroCluster configuration requires at least two SAS shelves. However, two shelves are recommended at each site to allow disk ownership on a per-shelf basis, for a total of four SAS shelves.

A minimum of one shelf at each site is supported.

Example names:

- Site A:
 - *shelf_A_1_1*
 - *shelf_A_1_2*
- Site B:
 - *shelf_B_1_1*
 - *shelf_B_1_2*

Information gathering worksheet for FC-to-SAS bridges

Before beginning to configure the MetroCluster sites, you should gather required configuration information.

Site A, FC-to-SAS bridge 1 (FC_bridge_A_1a)

Each SAS stack requires at least two FC-to-SAS bridges.

Each bridge connects to *Controller_A_1_port-number* and *Controller_B_1_port-number*.

Site A	Your value
Bridge_A_1a IP address	
Bridge_A_1a Username	
Bridge_A_1a Password	

Site A, FC-to-SAS bridge 2 (FC_bridge_A_1b)

Each SAS stack requires at least two FC-to-SAS bridges.

Each bridge connects to Controller_A_1_`port-number` and Controller_B_1_`port-number`.

Site A	Your value
Bridge_A_1b IP address	
Bridge_A_1b Username	
Bridge_A_1b Password	

Site B, FC-to-SAS bridge 1 (FC_bridge_B_1a)

Each SAS stack requires at least two FC-to-SAS bridges.

Each bridge connects to Controller_A_1_`port-number` and Controller_B_1_`port-number`.

Site B	Your value
Bridge_B_1a IP address	
Bridge_B_1a Username	
Bridge_B_1a Password	

Site B, FC-to-SAS bridge 2 (FC_bridge_B_1b)

Each SAS stack requires at least two FC-to-SAS bridges.

Each bridge connects to Controller_A_1_`port-number` and Controller_B_1_`port-number`.

Site B	Your value
Bridge_B_1b IP address	
Bridge_B_1b Username	

Install and cable MetroCluster components

Racking the hardware components

If you have not received the equipment already installed in cabinets, you must rack the components.

This task must be performed on both MetroCluster sites.

Steps

1. Plan out the positioning of the MetroCluster components.

The rack space depends on the platform model of the storage controllers, switch types, and the number of disk shelf stacks in your configuration.

2. Properly ground yourself.
3. Install the storage controllers in the rack or cabinet.

[ONTAP Hardware Systems Documentation](#)

4. Install the disk shelves, power them on, and set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites).
5. Install each FC-to-SAS bridge:
 - a. Secure the “L” brackets on the front of the bridge to the front of the rack (flush-mount) with the four screws.

The openings in the bridge “L” brackets are compliant with rack standard ETA-310-X for 19-inch (482.6 mm) racks.

For more information and an illustration of the installation, see the *ATTO FibreBridge Installation and Operation Manual for your bridge model*.

- b. Connect each bridge to a power source that provides a proper ground.
- c. Power on each bridge.



For maximum resiliency, bridges that are attached to the same stack of disk shelves must be connected to different power sources.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed its power-on self test sequence.

Cabling the controllers to each other

Each controller’s FC-VI adapters must be cabled directly to its partner.

Steps

1. Cable the FC-VI ports.



The above illustration is a typical representation of the required cabling. The specific FC-VI ports vary by controller module.

- AFF A300 and FAS8200 controller modules can be ordered with one of two options for FC-VI connectivity:
 - Onboard ports 0e and 0f configured in FC-VI mode.
 - Ports 1a and 1b on an FC-VI card in slot 1.
- AFF A700 and FAS9000 storage systems controller modules use four FC-VI ports each.

Cabling the cluster peering connections

You must cable the controller module ports used for cluster peering so that they have connectivity with the cluster on their partner site.

This task must be performed on each controller module in the MetroCluster configuration.

At least two ports on each controller module should be used for cluster peering.

The recommended minimum bandwidth for the ports and network connectivity is 1 GbE.

Steps

1. Identify and cable at least two ports for cluster peering and verify they have network connectivity with the partner cluster.

Cluster peering can be done on dedicated ports or on data ports. Using dedicated ports provides a higher throughput for the cluster peering traffic.

[Cluster and SVM peering express configuration](#)

Cabling the management and data connections

You must cable the management and data ports on each storage controller to the site networks.

This task must be repeated for each new controller at both MetroCluster sites.

You can connect the controller and cluster switch management ports to existing switches in your network. In addition you can connect controller to new dedicated network switches such as NetApp CN1601 cluster management switches.

Steps

1. Cable the controller's management and data ports to the management and data networks at the local site.

Install FC-to-SAS bridges and SAS disk shelves

Install and cable ATTO FibreBridge bridges and SAS disk shelves when you add new storage to the configuration.

About this task

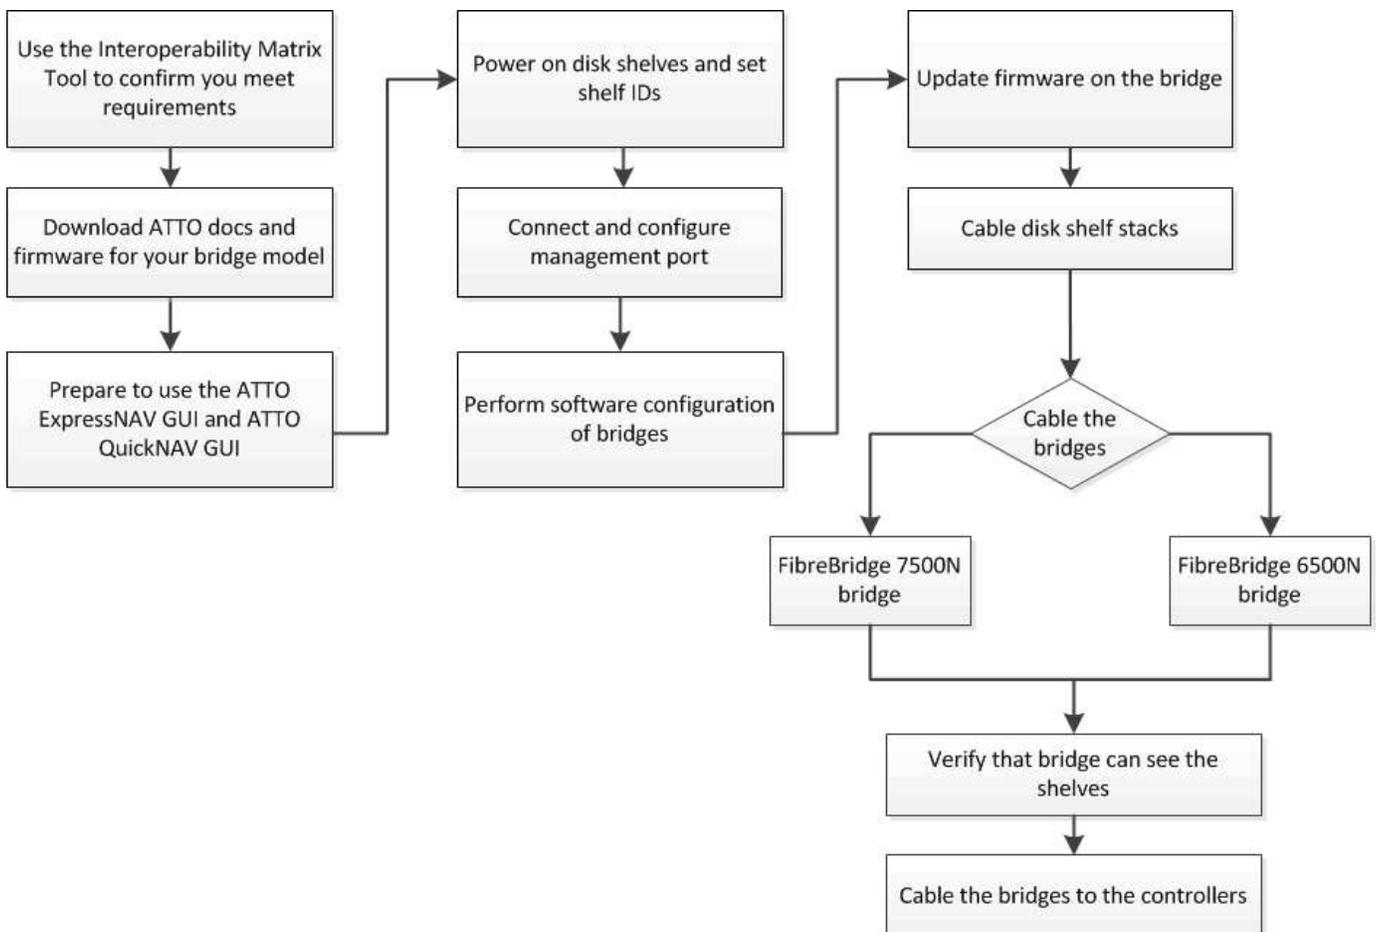
For systems received from the factory, the FC-to-SAS bridges are preconfigured and do not require additional configuration.

This procedure is written with the assumption that you are using the recommended bridge management interfaces: the ATTO ExpressNAV GUI and ATTO QuickNAV utility.

You use the ATTO ExpressNAV GUI to configure and manage a bridge, and to update the bridge firmware. You use the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port.

You can use other management interfaces instead, if needed, such as a serial port or Telnet to configure and manage a bridge and to configure the Ethernet management 1 port, and FTP to update the bridge firmware.

This procedure uses the following workflow:



In-band management of the FC-to-SAS bridges

Beginning with ONTAP 9.5 with FibreBridge 7500N or 7600N bridges, *in-band management* of the bridges is

supported as an alternative to IP management of the bridges. Beginning with ONTAP 9.8, out-of-band management is deprecated.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

When using in-band management, the bridges can be managed and monitored from the ONTAP CLI using the FC connection to the bridge. Physical access to the bridge through the bridge Ethernet ports is not required, reducing the security vulnerability of the bridge.

The availability of in-band management of the bridges depends on the version of ONTAP:

- Beginning with ONTAP 9.8, bridges are managed via in-band connections by default and out-of-band management of the bridges via SNMP is deprecated.
- ONTAP 9.5 through 9.7: Either in-band management or out-of-band SNMP management is supported.
- Before ONTAP 9.5, only out-of-band SNMP management is supported.

Bridge CLI commands can be issued from the ONTAP interface `storage bridge run-cli -name <bridge_name> -command <bridge_command_name> command` at the ONTAP interface.



Using in-band management with IP access disabled is recommended to improve security by limiting physical connectivity the bridge.

FibreBridge 7600N and 7500N bridge limits and attachment rules

Review the limits and considerations when attaching FibreBridge 7600N and 7500N bridges.

FibreBridge 7600N and 7500N bridge limits

- The maximum number of HDD and SSD drives combined is 240.
- The maximum number of SSD drives is 96.
- The maximum number of SSDs per SAS port is 48.
- The maximum number of shelves per SAS port is 10.

FibreBridge 7600N and 7500N bridge attachment rules

- Do not mix SSD and HDD drives on the same SAS port.
- Distribute the shelves evenly across the SAS ports.
- You shouldn't have DS460 shelves on the same SAS port as other shelf types (for example, DS212 or DS224 shelves).

Example configuration

The following shows an example configuration for connecting four DS224 shelves with SSD drives and six DS224 shelves with HDD drives:

SAS port	Shelves and drives
SAS port-A	2x DS224 shelves with SSD drives
SAS port-B	2x DS224 shelves with SSD drives

SAS port	Shelves and drives
SAS port-C	3x DS224 shelves with HDD drives
SAS port-D	3x DS224 shelves with HDD drives

Prepare for the installation

When you are preparing to install the bridges as part of your new MetroCluster system, you must verify that your system meets certain requirements, including meeting setup and configuration requirements for the bridges. Other requirements include downloading the necessary documents, the ATTO QuickNAV utility, and the bridge firmware.

Before you begin

- Your system must already be installed in a rack if it was not shipped in a system cabinet.
- Your configuration must be using supported hardware models and software versions.

In the [NetApp Interoperability Matrix Tool \(IMT\)](#), you can use the **Storage Solution** field to select your MetroCluster solution. You can use the **Component Explorer** to select the components and ONTAP version to refine your search. You can select **Show Results** to display the list of supported configurations that match the criteria.

- Each FC controller must have one FC port available for one bridge to connect to it.
- You must be familiar with how to handle SAS cables and the considerations and best practices for installing and cabling disk shelves.

The *Installation and Service Guide* for your disk shelf model describes the considerations and best practices.

- The computer you are using to set up the bridges must be running an ATTO-supported web browser to use the ATTO ExpressNAV GUI.

The *ATTO Product Release Notes* have an up-to-date list of supported web browsers. You can access this document from the ATTO web site as described in the following steps.

Steps

1. Download the *Installation and Service Guide* for your disk shelf model:
 - a. Access the ATTO web site using the link provided for your FibreBridge model and download the manual and the QuickNAV utility.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this and other content on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

2. Gather the hardware and information needed to use the recommended bridge management interfaces, the ATTO ExpressNAV GUI, and the ATTO QuickNAV utility:
 - a. Determine a non-default user name and password (for accessing the bridges).

You should change the default user name and password.

- b. If configuring for IP management of the bridges, you need the shielded Ethernet cable provided with the bridges (which connects from the bridge Ethernet management 1 port to your network).
- c. If configuring for IP management of the bridges, you need an IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge.
- d. Disable VPN clients on the computer you are using for setup.

Active VPN clients cause the QuickNAV scan for bridges to fail.

Install the FC-to-SAS bridge and SAS shelves

After ensuring that the system meets all of the requirements in “Preparing for the installation”, you can install your new system.

About this task

- The disk and shelf configuration at both sites should be identical.

If a non-mirrored aggregate is used, the disk and shelf configuration at each site might be different.



All disks in the disaster recovery group must use the same type of connection and be visible to all of the nodes within the disaster recovery group, regardless of the disks being used for mirrored or non-mirrored aggregate.

- The system connectivity requirements for maximum distances for disk shelves, FC controllers, and backup tape devices using 50-micron, multimode fiber-optic cables, also apply to FibreBridge bridges.

[NetApp Hardware Universe](#)



In-band ACP is supported without additional cabling in the following shelves and FibreBridge 7500N or 7600N bridge:

- IOM12 (DS460C) behind a 7500N or 7600N bridge with ONTAP 9.2 and later
- IOM12 (DS212C and DS224C) behind a 7500N or 7600N bridge with ONTAP 9.1 and later



SAS shelves in MetroCluster configurations do not support ACP cabling.

Enable IP port access on the FibreBridge 7600N bridge if necessary

If you are using an ONTAP version prior to 9.5, or otherwise plan to use out-of-band access to the FibreBridge 7600N bridge using telnet or other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV), you can enable the access services via the console port.

About this task

Unlike the ATTO FibreBridge 7500N bridges, the FibreBridge 7600N bridge is shipped with all IP port protocols and services disabled.

Beginning with ONTAP 9.5, *in-band management* of the bridges is supported. This means the bridges can be configured and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required and the bridge user interfaces are not required.

Beginning with ONTAP 9.8, *in-band management* of the bridges is supported by default and out-of-band SNMP management is deprecated.

This task is required if you are **not** using in-band management to manage the bridges. In this case, you need to configure the bridge via the Ethernet management port.

Steps

1. Access the bridge console interface by connecting a serial cable to the serial port on the FibreBridge 7600N bridge.

2. Using the console, enable the access services, and then save the configuration:

```
set closeport none
```

```
saveconfiguration
```

The `set closeport none` command enables all access services on the bridge.

3. Disable a service, if desired, by issuing the `set closeport` command and repeating the command as necessary until all desired services are disabled:

```
set closeport service
```

The `set closeport` command disables a single service at a time.

The parameter *service* can be specified as one of the following:

- `expressnav`
- `ftp`
- `icmp`
- `quicknav`
- `snmp`
- `telnet`

You can check whether a specific protocol is enabled or disabled by using the `get closeport` command.

4. If you are enabling SNMP, you must also issue following command:

```
set SNMP enabled
```

SNMP is the only protocol that requires a separate enable command.

5. Save the configuration:

```
saveconfiguration
```

Configure the FC-to-SAS bridges

Before cabling your model of the FC-to-SAS bridges, you must configure the settings in the FibreBridge software.

Before you begin

You should decide whether you will be using in-band management of the bridges.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.

Steps

1. Configure the serial console port on the ATTO FibreBridge by setting the port speed to 115000 bauds:

```
get serialportbaudrate
SerialPortBaudRate = 115200

Ready.

set serialportbaudrate 115200

Ready. *
saveconfiguration
Restart is necessary....
Do you wish to restart (y/n) ? y
```

2. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

3. If configuring for IP management, connect the Ethernet management 1 port on each bridge to your network by using an Ethernet cable.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

The Ethernet management 1 port enables you to quickly download the bridge firmware (using ATTO ExpressNAV or FTP management interfaces) and to retrieve core files and extract logs.

4. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

5. Configure the bridge.

You should make note of the user name and password that you designate.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

a. If configuring for IP management, configure the IP settings of the bridge.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

To set the IP address without the QuickNAV utility, you need to have a serial connection to the FibreBridge.

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address  
  
set ipsubnetmask mp1 subnet-mask  
  
set ipgateway mp1 x.x.x.x  
  
set ipdhcp mp1 disabled  
  
set ethernetspeed mp1 1000
```

b. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename <bridge_name>
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

6. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the FC port of the controller module to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate <port-number> <port-speed>
```

b. If you are configuring a FibreBridge 7500N bridge, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode <port-number> ptp
```

c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable <port-number>
```

The following example shows the disabling of FC port 2:

```
FCPortDisable 2
```

```
Fibre Channel Port 2 has been disabled.
```

d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used.

If only SAS port A is used, then SAS ports B, C, and D must be disabled. The following example shows the disabling of SAS port B. You must similarly disable SAS ports C and D:

```
SASPortDisable b
```

```
SAS Port B has been disabled.
```

7. Secure access to the bridge and save the bridge's configuration. Choose an option from below depending on the version of ONTAP your system is running.

ONTAP version	Steps
ONTAP 9.5 or later	<p>a. View the status of the bridges:</p> <pre>storage bridge show</pre> <p>The output shows which bridge is not secured.</p> <p>b. Secure the bridge:</p> <pre>securebridge</pre>
ONTAP 9.4 or earlier	<p>a. View the status of the bridges:</p> <pre>storage bridge show</pre> <p>The output shows which bridge is not secured.</p> <p>b. Check the status of the unsecured bridge's ports:</p> <pre>info</pre> <p>The output shows the status of Ethernet ports MP1 and MP2.</p> <p>c. If Ethernet port MP1 is enabled, run:</p> <pre>set EthernetPort mp1 disabled</pre> <p>If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.</p> <p>d. Save the bridge's configuration.</p> <p>You must run the following commands:</p> <pre>SaveConfiguration</pre> <pre>FirmwareRestart</pre> <p>You are prompted to restart the bridge.</p>

8. After completing MetroCluster configuration, use the `flashimages` command to check your version of FibreBridge firmware and, if the bridges are not using the latest supported version, update the firmware on

all bridges in the configuration.

Maintain MetroCluster Components

Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM12 modules

After configuring the bridge, you can start cabling your new system.

About this task

For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

Steps

1. Daisy-chain the disk shelves in each stack:
 - a. Beginning with the logical first shelf in the stack, connect IOM A port 3 to the to IOM A port 1 on the next shelf until each IOM A in the stack is connected.
 - b. Repeat the previous substep for IOM B.
 - c. Repeat the previous substeps for each stack.

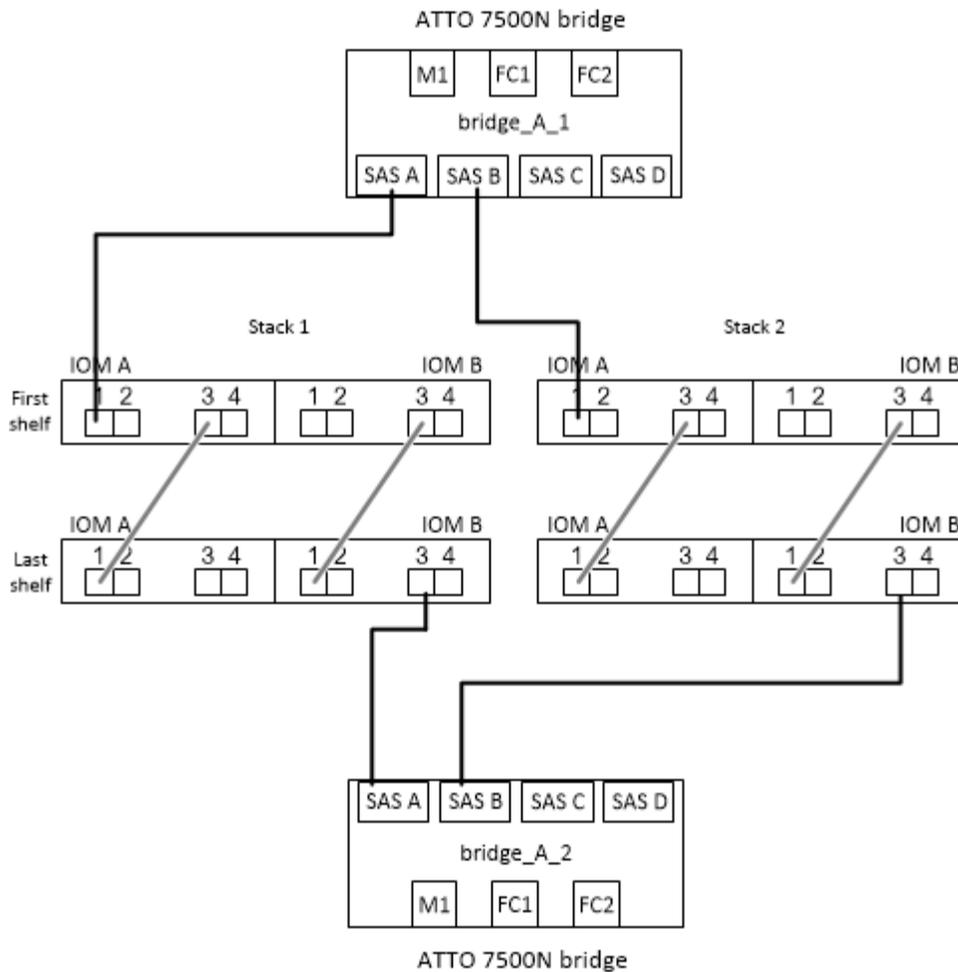
The *Installation and Service Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

2. Power on the disk shelves, and then set the shelf IDs.
 - You must power-cycle each disk shelf.
 - Shelf IDs must be unique for each SAS disk shelf within each MetroCluster DR group (including both sites).
3. Cable disk shelves to the FibreBridge bridges.
 - a. For the first stack of disk shelves, cable IOM A of the first shelf to SAS port A on FibreBridge A, and cable IOM B of the last shelf to SAS port A on FibreBridge B.
 - b. For additional shelf stacks, repeat the previous step using the next available SAS port on the FibreBridge bridges, using port B for the second stack, port C for the third stack, and port D for the fourth stack.
 - c. During cabling, attach the stacks based on IOM12 modules to the same bridge as long as they are connected to separate SAS ports.



Each stack can use different models of IOM, but all disk shelves within a stack must use the same model.

The following illustration shows disk shelves connected to a pair of FibreBridge 7600N or 7500N bridges:



Verify bridge connectivity and cable the FC-to-SAS bridges to the controller FC ports

You must cable the bridges to the controller FC ports in a two-node bridge-attached MetroCluster configuration.

Steps

1. Verify that each bridge can detect all of the disk drives and disk shelves to which the bridge is connected:

```
sastargets
```

The `sastargets` command output shows the devices (disks and disk shelves) connected to the bridge. The output lines are sequentially numbered so that you can quickly count the devices.

The following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

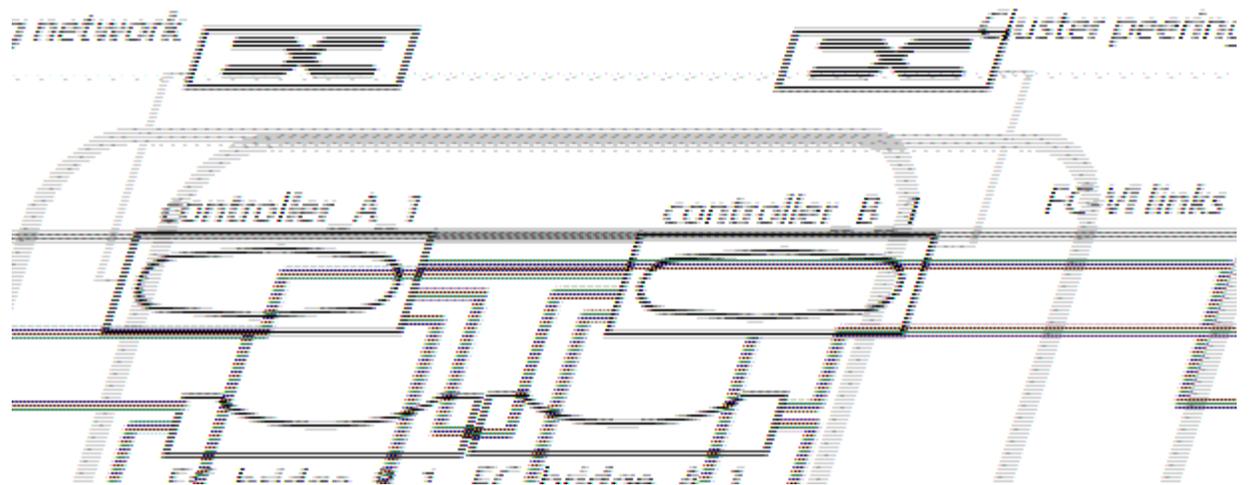
- Verify that the command output shows that the bridge is connected to the correct disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat Step 1 for each remaining bridge.
Not correct	<ol style="list-style-type: none"> Check for loose SAS cables or correct the SAS cabling by recabling the disk shelves to the bridges. Cable a FibreBridge 7600N or 7500N bridge with disk shelves using IOM12 modules Repeat Step 1 for each remaining bridge.

- Cable each bridge to the controller FC ports:
 - Cable FC port 1 of the bridge to an FC port on the controller in cluster_A.
 - Cable FC port 2 of the bridge to an FC port on the controller in cluster_B.
 - If the controller is configured with a quad-port FC adapter, make sure that the bridges at either end of the storage stack are not connected to two FC ports on the same ASIC. For example:
 - Port a and port b share the same ASIC.
 - Port c and port d share the same ASIC.

In this example, connect FC_bridge_A_1 to port a and FC_bridge_A2 to port c.
 - If the controller is configured with more than one FC adapter, do not cable the bridges at either end of the storage stack to the same adapter.

In this scenario, you should connect FC_bridge_A_1 to an onboard FC port, and connect FC_bridge_A_2 to an FC port on an adapter in an expansion slot.



4. Repeat [Step 3](#) on the other bridges until all of the bridges have been cabled.

Secure or unsecure the FibreBridge bridge

To easily disable potentially unsecure Ethernet protocols on a bridge, beginning with ONTAP 9.5 you can secure the bridge. This disables the bridge's Ethernet ports. You can also reenable Ethernet access.

About this task

- Securing the bridge disables telnet and other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV) on the bridge.
- This procedure uses out-of-band management using the ONTAP prompt, which is available beginning with ONTAP 9.5.

You can issue the commands from the bridge CLI if you are not using out-of-band management.

- The `unsecurebridge` command can be used to re-enable the Ethernet ports.
- In ONTAP 9.7 and earlier, running the `securebridge` command on the ATTO FibreBridge might not update the bridge status correctly on the partner cluster. If this occurs, run the `securebridge` command from the partner cluster.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. From the ONTAP prompt of the cluster containing the bridge, secure or unsecure the bridge.

- The following command secures `bridge_A_1`:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command securebridge
```

- The following command unsecures `bridge_A_1`:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command unsecurebridge
```

2. From the ONTAP prompt of the cluster containing the bridge, save the bridge configuration:

```
storage bridge run-cli -bridge <bridge-name> -command saveconfiguration
```

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command  
saveconfiguration
```

3. From the ONTAP prompt of the cluster containing the bridge, restart the bridge's firmware:

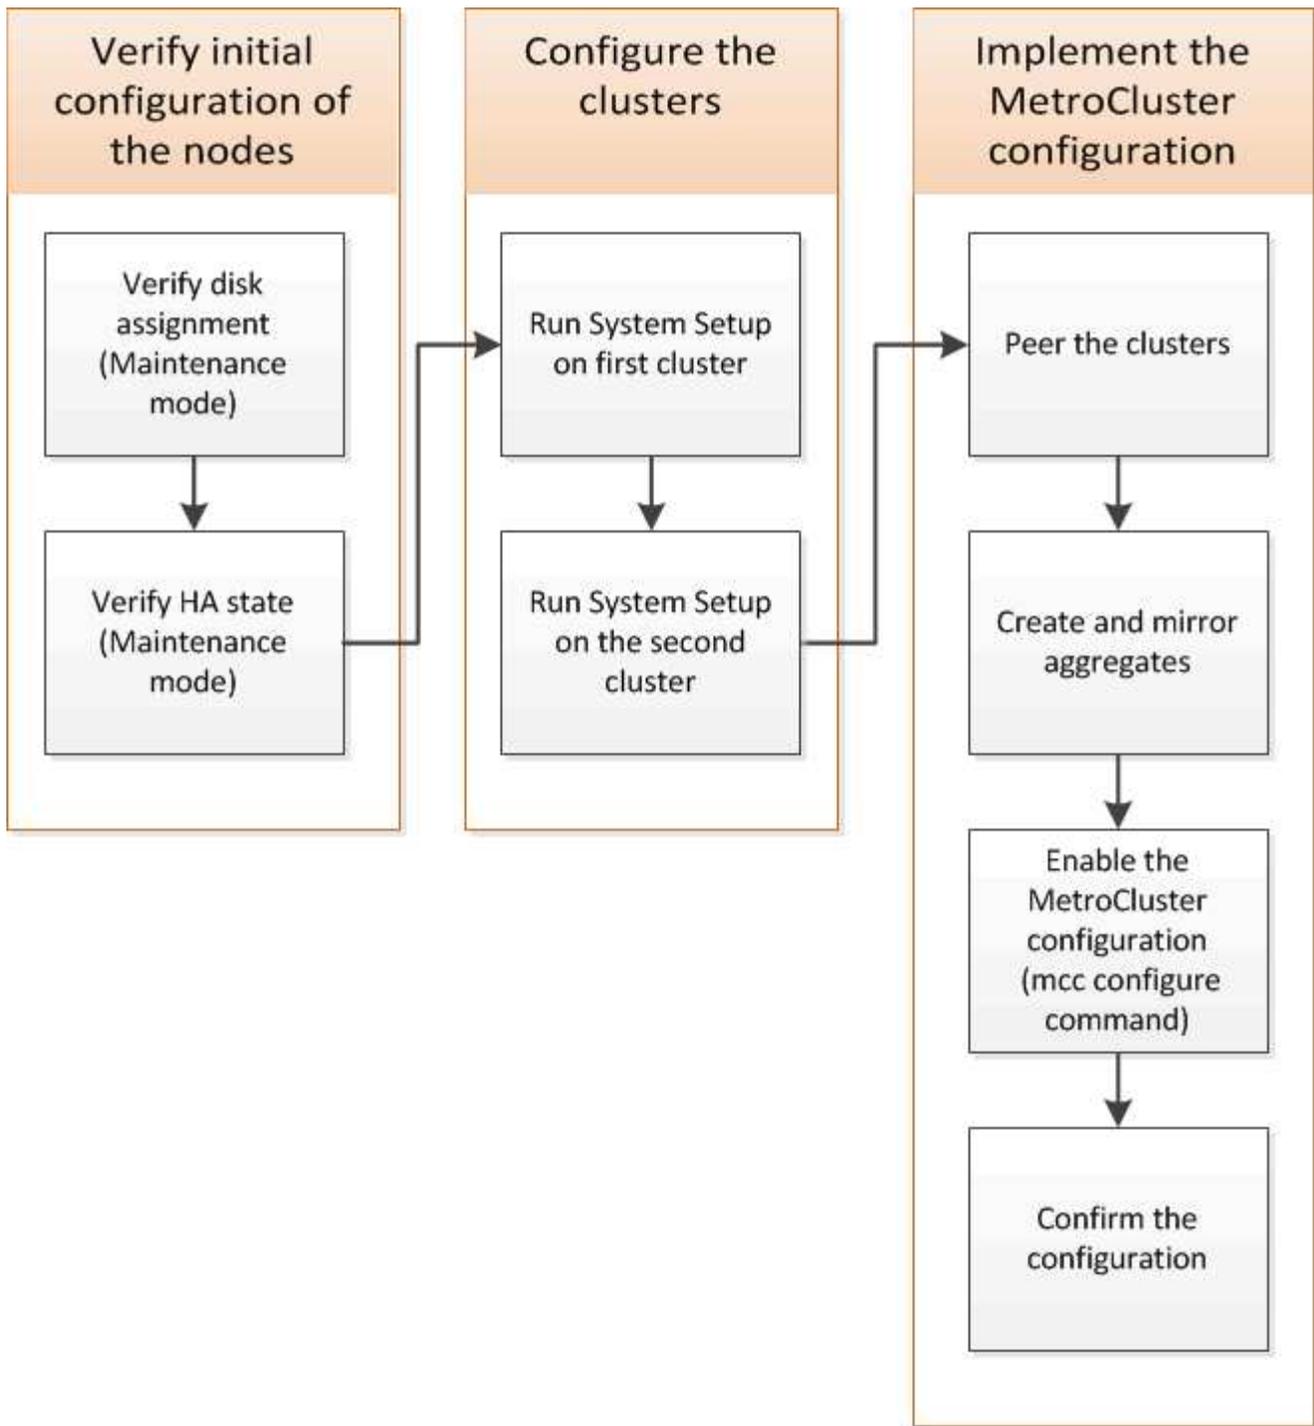
```
storage bridge run-cli -bridge <bridge-name> -command firmwarerestart
```

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command firmwarerestart
```

Configuring the MetroCluster software in ONTAP

You must set up each node in the MetroCluster configuration in ONTAP, including the node-level configurations and the configuration of the nodes into two sites. You must also implement the MetroCluster relationship between the two sites.



Steps

1. Gather the required IP addresses for the controller modules before you begin the configuration process.
2. Complete the IP network information worksheet for site A.

IP network information worksheet for Site A

You must obtain IP addresses and other network information for the first MetroCluster site (site A) from your network administrator before you configure the system.

Site A cluster creation information

When you first create the cluster, you need the following information:

Type of information	Your values
Cluster name. Example used in this information: site_A	
DNS domain	
DNS name servers	
Location	
Administrator password	

Site A node information

For each node in the cluster, you need a management IP address, a network mask, and a default gateway.

Node	Port	IP address	Network mask	Default gateway
Node 1. Example used in this information: controller_A_1				
Node 2. Not required if using two-node MetroCluster configuration (one node at each site). Example used in this information: controller_A_2				

Site A LIFs and ports for cluster peering

For each node in the cluster, you need the IP addresses of two intercluster LIFs, including a network mask and a default gateway. The intercluster LIFs are used to peer the clusters.

Node	Port	IP address of intercluster LIF	Network mask	Default gateway
Node 1 IC LIF 1				
Node 1 IC LIF 2				

Site A time server information

You must synchronize the time, which requires one or more NTP time servers.

Node	Host name	IP address	Network mask	Default gateway
NTP server 1				
NTP server 2				

Site A AutoSupport information

You must configure AutoSupport on each node, which requires the following information:

Type of information	Your values	
From email address		
Mail hosts	IP addresses or names	
Transport protocol	HTTP, HTTPS, or SMTP	
	Proxy server	
Recipient email addresses or distribution lists	Full-length messages	
	Concise messages	
	Partners	

Site A SP information

You must enable access to the Service Processor (SP) of each node for troubleshooting and maintenance. This requires the following network information for each node:

Node	IP address	Network mask	Default gateway
Node 1			

IP network information worksheet for site B

You must obtain IP addresses and other network information for the second MetroCluster site (site B) from your network administrator before you configure the system.

Site B cluster creation information

When you first create the cluster, you need the following information:

Type of information	Your values
Cluster name. Example used in this information: site_B	
DNS domain	
DNS name servers	
Location	
Administrator password	

Site B node information

For each node in the cluster, you need a management IP address, a network mask, and a default gateway.

Node	Port	IP address	Network mask	Default gateway
Node 1. Example used in this information: controller_B_1				
Node 2. Not required for two-node MetroCluster configurations (one node at each site). Example used in this information: controller_B_2				

Site B LIFs and ports for cluster peering

For each node in the cluster, you need the IP addresses of two intercluster LIFs, including a network mask and a default gateway. The intercluster LIFs are used to peer the clusters.

Node	Port	IP address of intercluster LIF	Network mask	Default gateway
Node 1 IC LIF 1				
Node 1 IC LIF 2				

Site B time server information

You must synchronize the time, which requires one or more NTP time servers.

Node	Host name	IP address	Network mask	Default gateway
NTP server 1				
NTP server 2				

Site B AutoSupport information

You must configure AutoSupport on each node, which requires the following information:

Type of information	Your values	
From email address		
Mail hosts	IP addresses or names	
Transport protocol	HTTP, HTTPS, or SMTP	
	Proxy server	
Recipient email addresses or distribution lists	Full-length messages	
	Concise messages	
	Partners	

Site B SP information

You must enable access to the Service Processor (SP) of each node for troubleshooting and maintenance, which requires the following network information for each node:

Node	IP address	Network mask	Default gateway
Node 1 (controller_B_1)			

Similarities and differences between standard cluster and MetroCluster configurations

The configuration of the nodes in each cluster in a MetroCluster configuration is similar to that of nodes in a standard cluster.

The MetroCluster configuration is built on two standard clusters. Physically, the configuration must be symmetrical, with each node having the same hardware configuration, and all of the MetroCluster components must be cabled and configured. However, the basic software configuration for nodes in a MetroCluster configuration is the same as that for nodes in a standard cluster.

Configuration step	Standard cluster configuration	MetroCluster configuration
--------------------	--------------------------------	----------------------------

Configure management, cluster, and data LIFs on each node.	Same in both types of clusters	
Configure the root aggregate.	Same in both types of clusters	
Set up the cluster on one node in the cluster.	Same in both types of clusters	
Join the other node to the cluster.	Same in both types of clusters	
Create a mirrored root aggregate.	Optional	Required
Peer the clusters.	Optional	Required
Enable the MetroCluster configuration.	Does not apply	Required

Restoring system defaults and configuring the HBA type on a controller module

To ensure a successful MetroCluster installation, reset and restore defaults on the controller modules.

Important

This task is only required for stretch configurations using FC-to-SAS bridges.

Steps

1. At the LOADER prompt, return the environmental variables to their default setting:

```
set-defaults
```

2. Boot the node into Maintenance mode, then configure the settings for any HBAs in the system:

- a. Boot into Maintenance mode:

```
boot_ontap maint
```

- b. Check the current settings of the ports:

```
ucadmin show
```

- c. Update the port settings as needed.

If you have this type of HBA and desired mode...	Use this command...
CNA FC	<code>ucadmin modify -m fc -t initiator <i>adapter_name</i></code>
CNA Ethernet	<code>ucadmin modify -mode cna <i>adapter_name</i></code>

FC target	<code>fcadmin config -t target <i>adapter_name</i></code>
FC initiator	<code>fcadmin config -t initiator <i>adapter_name</i></code>

3. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

4. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

```
boot_ontap maint
```

5. Verify the changes you made:

If you have this type of HBA...	Use this command...
CNA	<code>ucadmin show</code>
FC	<code>fcadmin show</code>

6. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

7. Boot the node to the boot menu:

```
boot_ontap menu
```

After you run the command, wait until the boot menu is shown.

8. Clear the node configuration by typing “wipeconfig” at the boot menu prompt, and then press Enter.

The following screen shows the boot menu prompt:

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning.

Selection (1-9)? wipeconfig

This option deletes critical system configuration, including cluster membership.

Warning: do not run this option on a HA node that has been taken over.

Are you sure you want to continue?: yes

Rebooting to finish wipeconfig request.

Configuring FC-VI ports on a X1132A-R6 quad-port card on FAS8020 systems

If you are using the X1132A-R6 quad-port card on a FAS8020 system, you can enter Maintenance mode to configure the 1a and 1b ports for FC-VI and initiator usage. This is not required on MetroCluster systems received from the factory, in which the ports are set appropriately for your configuration.

About this task

This task must be performed in Maintenance mode.



Converting an FC port to an FC-VI port with the `ucadmin` command is only supported on the FAS8020 and AFF 8020 systems. Converting FC ports to FCVI ports is not supported on any other platform.

Steps

1. Disable the ports:

```
storage disable adapter 1a
```

```
storage disable adapter 1b
```

```
*> storage disable adapter 1a
Jun 03 02:17:57 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1a.
Host adapter 1a disable succeeded
Jun 03 02:17:57 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1a is now offline.
*> storage disable adapter 1b
Jun 03 02:18:43 [controller_B_1:fc.adapter.offlining:info]: Offlining
Fibre Channel adapter 1b.
Host adapter 1b disable succeeded
Jun 03 02:18:43 [controller_B_1:fc.adapter.offline:info]: Fibre Channel
adapter 1b is now offline.
*>
```

2. Verify that the ports are disabled:

```
ucadmin show
```

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
...					
1a	fc	initiator	-	-	offline
1b	fc	initiator	-	-	offline
1c	fc	initiator	-	-	online
1d	fc	initiator	-	-	online

3. Set the a and b ports to FC-VI mode:

```
ucadmin modify -adapter 1a -type fcvi
```

The command sets the mode on both ports in the port pair, 1a and 1b (even though only 1a is specified in the command).

```
*> ucadmin modify -t fcvi 1a
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1a. Reboot the controller for the changes to
take effect.
Jun 03 02:19:13 [controller_B_1:ucm.type.changed:info]: FC-4 type has
changed to fcvi on adapter 1b. Reboot the controller for the changes to
take effect.
```

4. Confirm that the change is pending:

```
ucadmin show
```

```
*> ucadmin show
      Current   Current   Pending   Pending   Admin
Adapter Mode     Type     Mode     Type     Status
-----
...
1a    fc      initiator -        fcvi    offline
1b    fc      initiator -        fcvi    offline
1c    fc      initiator -         -      online
1d    fc      initiator -         -      online
```

- 5. Shut down the controller, and then reboot into Maintenance mode.
- 6. Confirm the configuration change:

```
ucadmin show local
```

```
Node           Adapter  Mode     Type           Mode     Type           Status
-----
...
controller_B_1 1a       fc       fcvi           -        -              online
controller_B_1 1b       fc       fcvi           -        -              online
controller_B_1 1c       fc       initiator      -        -              online
controller_B_1 1d       fc       initiator      -        -              online
6 entries were displayed.
```

Verifying disk assignment in Maintenance mode in a two-node configuration

Before fully booting the system to ONTAP, you can optionally boot the system to Maintenance mode and verify the disk assignment on the nodes. The disks should be assigned to create a fully symmetric configuration with both sites owning their own disk shelves and serving data, where each node and each pool have an equal number of mirrored disks assigned to them.

Before you begin

The system must be in Maintenance mode.

About this task

New MetroCluster systems have disk assignments completed prior to shipment.

The following table shows example pool assignments for a MetroCluster configuration. Disks are assigned to pools on a per-shelf basis.

Disk shelf (<i>example name</i>)...	At site...	Belongs to...	And is assigned to that node's...
Disk shelf 1 (shelf_A_1_1)	Site A	Node A 1	Pool 0
Disk shelf 2 (shelf_A_1_3)			
Disk shelf 3 (shelf_B_1_1)		Node B 1	Pool 1
Disk shelf 4 (shelf_B_1_3)			
Disk shelf 9 (shelf_B_1_2)	Site B	Node B 1	Pool 0
Disk shelf 10 (shelf_B_1_4)			
Disk shelf 11 (shelf_A_1_2)		Node A 1	Pool 1
Disk shelf 12 (shelf_A_1_4)			

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

Assign these disks in the drawer...	To this node and pool...
1 - 6	Local node's pool 0
7 - 12	DR partner's pool 1

This disk assignment pattern minimizes the effect on an aggregate if a drawer goes offline.

Steps

1. If your system was received from the factory, confirm the shelf assignments:

```
disk show -v
```

2. If necessary, you can explicitly assign disks on the attached disk shelves to the appropriate pool

```
disk assign
```

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1. You should assign an equal number of shelves to each pool.

- a. If you have not done so, boot each system into Maintenance mode.
- b. On the node on site A, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf disk_shelf_name -p pool
```

If storage controller node_A_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf shelf_A_1_1 -p 0
*> disk assign -shelf shelf_A_1_3 -p 0

*> disk assign -shelf shelf_A_1_2 -p 1
*> disk assign -shelf shelf_A_1_4 -p 1
```

- c. On the node at the remote site (site B), systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf disk_shelf_name -p pool
```

If storage controller node_B_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf shelf_B_1_2 -p 0
*> disk assign -shelf shelf_B_1_4 -p 0

*> disk assign -shelf shelf_B_1_1 -p 1
*> disk assign -shelf shelf_B_1_3 -p 1
```

- d. Show the disk shelf IDs and bays for each disk:

```
disk show -v
```

Verifying the HA state of components

In a stretch MetroCluster configuration that is not preconfigured at the factory, you must verify that the HA state of the controller and chassis component is set to “mcc-2n” so that they boot up properly. For systems received from the factory, this value is preconfigured and you do not need to verify it.

Before you begin

The system must be in Maintenance mode.

Steps

1. In Maintenance mode, view the HA state of the controller module and chassis:

```
ha-config show
```

The controller module and chassis should show the value “mcc-2n”.

2. If the displayed system state of the controller is not “mcc-2n”, set the HA state for the controller:

```
ha-config modify controller mcc-2n
```

3. If the displayed system state of the chassis is not “mcc-2n”, set the HA state for the chassis:

```
ha-config modify chassis mcc-2n
```

Halt the node.

Wait until the node is back at the LOADER prompt.

4. Repeat these steps on each node in the MetroCluster configuration.

Setting up ONTAP in a two-node MetroCluster configuration

In a two-node MetroCluster configuration, on each cluster you must boot up the node, exit the Cluster Setup wizard, and use the `cluster setup` command to configure the node into a single-node cluster.

Before you begin

You must not have configured the Service Processor.

About this task

This task is for two-node MetroCluster configurations using native NetApp storage.

This task must be performed on both clusters in the MetroCluster configuration.

For more general information about setting up ONTAP, see the [Setup ONTAP](#)

Steps

1. Power on the first node.



You must repeat this step on the node at the disaster recovery (DR) site.

The node boots, then the Cluster Setup wizard starts on the console informing you that AutoSupport will be enabled automatically.

```
::> Welcome to the cluster setup wizard.
```

You can enter the following commands at any time:

```
"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.
```

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
resolution, should a problem occur on your system.
For further information on AutoSupport, see:
<http://support.netapp.com/autosupport/>

```
Type yes to confirm and continue {yes}: yes
```

```
Enter the node management interface port [e0M]:
```

```
Enter the node management interface IP address [10.101.01.01]:
```

```
Enter the node management interface netmask [101.010.101.0]:
```

```
Enter the node management interface default gateway [10.101.01.0]:
```

```
Do you want to create a new cluster or join an existing cluster?  
{create, join}:
```

2. Create a new cluster:

```
create
```

3. Choose whether the node is to be used as a single node cluster.

```
Do you intend for this node to be used as a single node cluster? {yes,  
no} [yes]:
```

4. Accept the system default "yes" by pressing Enter, or enter your own values by typing "no", and then

pressing Enter.

5. Follow the prompts to complete the **Cluster Setup** wizard, pressing Enter to accept the default values or typing your own values and then pressing Enter.

The default values are determined automatically based on your platform and network configuration.

6. After you complete the **Cluster Setup** wizard and it exits, verify that the cluster is active and the first node is healthy:

```
cluster show
```

The following example shows a cluster in which the first node (cluster1-01) is healthy and eligible to participate:

```
cluster1::> cluster show
Node                Health  Eligibility
-----
cluster1-01        true   true
```

If it becomes necessary to change any of the settings you entered for the admin SVM or node SVM, you can access the **Cluster Setup** wizard by using the `cluster setup` command.

Configuring the clusters into a MetroCluster configuration

You must peer the clusters, mirror the root aggregates, create a mirrored data aggregate, and then issue the command to implement the MetroCluster operations.

Peering the clusters

The clusters in the MetroCluster configuration must be in a peer relationship so that they can communicate with each other and perform the data mirroring essential to MetroCluster disaster recovery.

Related information

[Cluster and SVM peering express configuration](#)

[Considerations when using dedicated ports](#)

[Considerations when sharing data ports](#)

Configuring intercluster LIFs

You must create intercluster LIFs on ports used for communication between the MetroCluster partner clusters. You can use dedicated ports or ports that also have data traffic.

Configuring intercluster LIFs on dedicated ports

You can configure intercluster LIFs on dedicated ports. Doing so typically increases the available bandwidth for replication traffic.

Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in “cluster01”:

```
cluster01::> network port show
```

(Mbps)							Speed
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	

cluster01-01							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	
	e0e	Default	Default	up	1500	auto/1000	
	e0f	Default	Default	up	1500	auto/1000	
cluster01-02							
	e0a	Cluster	Cluster	up	1500	auto/1000	
	e0b	Cluster	Cluster	up	1500	auto/1000	
	e0c	Default	Default	up	1500	auto/1000	
	e0d	Default	Default	up	1500	auto/1000	
	e0e	Default	Default	up	1500	auto/1000	
	e0f	Default	Default	up	1500	auto/1000	

2. Determine which ports are available to dedicate to intercluster communication:

```
network interface show -fields home-port,curr-port
```

For complete command syntax, see the man page.

The following example shows that ports “e0e” and “e0f” have not been assigned LIFs:

```

cluster01::> network interface show -fields home-port,curr-port
vserver lif                home-port curr-port
Cluster cluster01-01_clus1 e0a        e0a
Cluster cluster01-01_clus2 e0b        e0b
Cluster cluster01-02_clus1 e0a        e0a
Cluster cluster01-02_clus2 e0b        e0b
cluster01
      cluster_mgmt         e0c        e0c
cluster01
      cluster01-01_mgmt1   e0c        e0c
cluster01
      cluster01-02_mgmt1   e0c        e0c

```

3. Create a failover group for the dedicated ports:

```

network interface failover-groups create -vserver system_SVM -failover-group
failover_group -targets physical_or_logical_ports

```

The following example assigns ports “e0e” and “e0f” to the failover group “intercluster01” on system SVM “cluster01”:

```

cluster01::> network interface failover-groups create -vserver cluster01
-failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f

```

4. Verify that the failover group was created:

```

network interface failover-groups show

```

For complete command syntax, see the man page.

```

cluster01::> network interface failover-groups show

```

Vserver	Group	Failover Targets
Cluster	Cluster	cluster01-01:e0a, cluster01-01:e0b, cluster01-02:e0a, cluster01-02:e0b
cluster01	Default	cluster01-01:e0c, cluster01-01:e0d, cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f
	intercluster01	cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

ONTAP version	Command
ONTAP 9.6 and later	network interface create -vserver system_SVM -lif LIF_name -service-policy default-intercluster -home -node node -home-port port -address port_IP -netmask netmask -failover-group failover_group
ONTAP 9.5 and earlier	network interface create -vserver system_SVM -lif LIF_name -role intercluster -home-node node -home-port port -address port_IP -netmask netmask -failover-group failover_group

For complete command syntax, see the man page.

The following example creates intercluster LIFs “cluster01_icl01” and “cluster01_icl02” in the failover group “intercluster01”:

```

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0 -failover-group intercluster01

```

6. Verify that the intercluster LIFs were created:

ONTAP version	Command
ONTAP 9.6 and later	network interface show -service-policy default-intercluster
ONTAP 9.5 and earlier	network interface show -role intercluster

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-intercluster
      Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
cluster01
      cluster01_icl01
              up/up      192.168.1.201/24  cluster01-01  e0e
true
      cluster01_icl02
              up/up      192.168.1.202/24  cluster01-02  e0f
true

```

7. Verify that the intercluster LIFs are redundant:

ONTAP version	Command
ONTAP 9.6 and later	network interface show -service-policy default-intercluster -failover

In ONTAP 9.5 and earlier

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs “cluster01_icl01” and “cluster01_icl02” on the SVM port “e0e” will fail over to port “e0f”.

```
cluster01::> network interface show -service-policy default-intercluster
-failover
          Logical          Home          Failover          Failover
Vserver  Interface          Node:Port      Policy            Group
-----  -
cluster01
          cluster01_icl01 cluster01-01:e0e  local-only
intercluster01
                                Failover Targets: cluster01-01:e0e,
                                                                cluster01-01:e0f
          cluster01_icl02 cluster01-02:e0e  local-only
intercluster01
                                Failover Targets: cluster01-02:e0e,
                                                                cluster01-02:e0f
```

Related information

[Considerations when using dedicated ports](#)

Configuring intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in “cluster01”:

```

cluster01::> network port show

(Mbps)
Node   Port      IPspace      Broadcast Domain Link   MTU   Admin/Oper
-----
cluster01-01
      e0a      Cluster      Cluster      up    1500  auto/1000
      e0b      Cluster      Cluster      up    1500  auto/1000
      e0c      Default     Default      up    1500  auto/1000
      e0d      Default     Default      up    1500  auto/1000
cluster01-02
      e0a      Cluster      Cluster      up    1500  auto/1000
      e0b      Cluster      Cluster      up    1500  auto/1000
      e0c      Default     Default      up    1500  auto/1000
      e0d      Default     Default      up    1500  auto/1000

```

2. Create intercluster LIFs on the system SVM:

ONTAP version	Command
ONTAP 9.6 and later	<code>network interface create -vserver <i>system_SVM</i> -lif <i>LIF_name</i> -service-policy default-intercluster -home -node <i>node</i> -home-port <i>port</i> -address <i>port_IP</i> -netmask <i>netmask</i></code>
ONTAP 9.5 and earlier	<code>network interface create -vserver <i>system_SVM</i> -lif <i>LIF_name</i> -role intercluster -home-node <i>node</i> -home-port <i>port</i> -address <i>port_IP</i> -netmask <i>netmask</i></code>

For complete command syntax, see the man page.

The following example creates intercluster LIFs “cluster01_icl01” and “cluster01_icl02”:

```

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0c
-address 192.168.1.201
-netmask 255.255.255.0

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0c
-address 192.168.1.202
-netmask 255.255.255.0

```

3. Verify that the intercluster LIFs were created:

ONTAP version	Command
ONTAP 9.6 and later	<code>network interface show -service-policy default-intercluster</code>
ONTAP 9.5 and earlier	<code>network interface show -role intercluster</code>

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-intercluster
          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node         Port
Home
-----
-----
cluster01
          cluster01_icl01
                up/up      192.168.1.201/24  cluster01-01  e0c
true
          cluster01_icl02
                up/up      192.168.1.202/24  cluster01-02  e0c
true

```

4. Verify that the intercluster LIFs are redundant:

ONTAP version	Command
ONTAP 9.6 and later	<code>network interface show -service-policy default-intercluster -failover</code>
ONTAP 9.5 and earlier	<code>network interface show -role intercluster -failover</code>

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs “cluster01_icl01” and “cluster01_icl02” on port “e0c” will fail over to port “e0d”.

```

cluster01::> network interface show -service-policy default-intercluster
-failover

```

Vserver	Logical Interface	Home Node:Port	Failover Policy	Failover Group
cluster01	cluster01_icl01	cluster01-01:e0c	local-only	
	192.168.1.201/24			
			Failover Targets: cluster01-01:e0c,	
			cluster01-01:e0d	
	cluster01_icl02	cluster01-02:e0c	local-only	
	192.168.1.201/24			
			Failover Targets: cluster01-02:e0c,	
			cluster01-02:e0d	

Related information

[Considerations when sharing data ports](#)

Creating a cluster peer relationship

You must create the cluster peer relationship between the MetroCluster clusters.

Creating a cluster peer relationship

You can use the `cluster peer create` command to create a peer relationship between a local and remote cluster. After the peer relationship has been created, you can run `cluster peer create` on the remote cluster to authenticate it to the local cluster.

Before you begin

- You must have created intercluster LIFs on every node in the clusters that are being peered.
- The clusters must be running ONTAP 9.3 or later.

Steps

1. On the destination cluster, create a peer relationship with the source cluster:

```

cluster peer create -generate-passphrase -offer-expiration MM/DD/YYYY
HH:MM:SS|1...7days|1...168hours -peer-addr peer_LIF_IPs -ip-space ip-space

```

If you specify both `-generate-passphrase` and `-peer-addr`, only the cluster whose intercluster LIFs are specified in `-peer-addr` can use the generated password.

You can ignore the `-ip-space` option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship on an unspecified remote cluster:

```
cluster02::> cluster peer create -generate-passphrase -offer-expiration
2days
```

```
                Passphrase: UCa+6lRVICXeL/gq1WrK7ShR
                Expiration Time: 6/7/2017 08:16:10 EST
Initial Allowed Vserver Peers: -
                Intercluster LIF IP: 192.140.112.101
                Peer Cluster Name: Clus_7ShR (temporary generated)
```

Warning: make a note of the passphrase - it cannot be displayed again.

2. On source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addr peer_LIF_IPs -ip-space ip-space
```

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses 192.140.112.101 and 192.140.112.102:

```
cluster01::> cluster peer create -peer-addr
192.140.112.101,192.140.112.102
```

Notice: Use a generated passphrase or choose a passphrase of 8 or more characters.

To ensure the authenticity of the peering relationship, use a phrase or sequence of characters that would be hard to guess.

```
Enter the passphrase:
Confirm the passphrase:
```

```
Clusters cluster02 and cluster01 are peered.
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

```

cluster01::> cluster peer show -instance

Peer Cluster Name: cluster02
Remote Intercluster Addresses: 192.140.112.101,
192.140.112.102
Availability of the Remote Cluster: Available
Remote Cluster Name: cluster2
Active IP Addresses: 192.140.112.101,
192.140.112.102

Cluster Serial Number: 1-80-123456
Address Family of Relationship: ipv4
Authentication Status Administrative: no-authentication
Authentication Status Operational: absent
Last Update Time: 02/05 21:05:41
IPspace for the Relationship: Default

```

4. Check the connectivity and status of the nodes in the peer relationship:

```
cluster peer health show
```

```

cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
-----
cluster01-01
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
cluster01-02
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true

```

Creating a cluster peer relationship (ONTAP 9.2 and earlier)

You can use the `cluster peer create` command to initiate a request for a peering relationship between a local and remote cluster. After the peer relationship has been requested by the local cluster, you can run

`cluster peer create` on the remote cluster to accept the relationship.

Before you begin

- You must have created intercluster LIFs on every node in the clusters being peered.
- The cluster administrators must have agreed on the passphrase each cluster will use to authenticate itself to the other.

Steps

1. On the data protection destination cluster, create a peer relationship with the data protection source cluster:

```
cluster peer create -peer-addr peer_LIF_IPs -ip-space ip-space
```

You can ignore the `-ip-space` option if you are not using a custom IPspace. For complete command syntax, see the man page.

The following example creates a cluster peer relationship with the remote cluster at intercluster LIF IP addresses 192.168.2.201 and 192.168.2.202:

```
cluster02::> cluster peer create -peer-addr 192.168.2.201,192.168.2.202
Enter the passphrase:
Please enter the passphrase again:
```

Enter the passphrase for the peer relationship when prompted.

2. On the data protection source cluster, authenticate the source cluster to the destination cluster:

```
cluster peer create -peer-addr peer_LIF_IPs -ip-space ip-space
```

For complete command syntax, see the man page.

The following example authenticates the local cluster to the remote cluster at intercluster LIF IP addresses 192.140.112.203 and 192.140.112.204:

```
cluster01::> cluster peer create -peer-addr 192.168.2.203,192.168.2.204
Please confirm the passphrase:
Please confirm the passphrase again:
```

Enter the passphrase for the peer relationship when prompted.

3. Verify that the cluster peer relationship was created:

```
cluster peer show -instance
```

For complete command syntax, see the man page.

```
cluster01::> cluster peer show -instance
Peer Cluster Name: cluster01
Remote Intercluster Addresses: 192.168.2.201,192.168.2.202
Availability: Available
Remote Cluster Name: cluster02
Active IP Addresses: 192.168.2.201,192.168.2.202
Cluster Serial Number: 1-80-000013
```

4. Check the connectivity and status of the nodes in the peer relationship:

```
cluster peer health show
```

For complete command syntax, see the man page.

```
cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
-----
cluster01-01
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
cluster01-02
          cluster02          cluster02-01
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          cluster02-02
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
```

Mirroring the root aggregates

You must mirror the root aggregates to provide data protection.

About this task

By default, the root aggregate is created as RAID-DP type aggregate. You can change the root aggregate from RAID-DP to RAID4 type aggregate. The following command modifies the root aggregate for RAID4 type aggregate:

```
storage aggregate modify -aggregate aggr_name -raidtype raid4
```



On non-ADP systems, the RAID type of the aggregate can be modified from the default RAID-DP to RAID4 before or after the aggregate is mirrored.

Steps

1. Mirror the root aggregate:

```
storage aggregate mirror aggr_name
```

The following command mirrors the root aggregate for “controller_A_1”:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote MetroCluster site.

2. Repeat the previous step for each node in the MetroCluster configuration.

Related information

[Logical storage management](#)

[ONTAP concepts](#)

Creating a mirrored data aggregate on each node

You must create a mirrored data aggregate on each node in the DR group.

Before you begin

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.

About this task

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.

[Disk and aggregate management](#)

- Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.

Steps

1. Display a list of available spares:

```
storage disk show -spare -owner node_name
```

2. Create the aggregate:

```
storage aggregate create -mirror true
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To ensure that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include



In the minimum supported configuration, in which a limited number of drives are available, you must use the `force-small-aggregate` option to allow the creation of a three disk RAID-DP aggregate.

- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create man` page.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10
-node node_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_node_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate aggregate-name
```

Creating unmirrored data aggregates

You can optionally create unmirrored data aggregates for data that does not require the redundant mirroring provided by MetroCluster configurations.

Before you begin

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can verify that the correct drive type is selected.

Example 1. About this task

ATTENTION: In MetroCluster FC configurations, the unmirrored aggregates will only be online after a switchover if the remote disks in the aggregate are accessible. If the ISLs fail, the local node may be unable to access the data in the unmirrored remote disks. The failure of an aggregate can lead to a reboot of the local node.



The unmirrored aggregates must be local to the node owning them.

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- The [Disks and aggregates management](#) contains more information about mirroring aggregates.

Steps

1. Display a list of available spares:

```
storage disk show -spare -owner node_name
```

2. Create the aggregate:

```
storage aggregate create
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To verify that the aggregate is created on a specific node, you should use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
 - List of specific drives that are to be added to the aggregate
 - Number of drives to include
 - Checksum style to use for the aggregate
 - Type of drives to use
 - Size of drives to use
 - Drive speed to use
 - RAID type for RAID groups on the aggregate
 - Maximum number of drives that can be included in a RAID group
 - Whether drives with different RPM are allowed
- For more information about these options, see the `storage aggregate create man` page.

The following command creates a unmirrored aggregate with 10 disks:

```

controller_A_1::> storage aggregate create aggr1_controller_A_1
-diskcount 10 -node controller_A_1
[Job 15] Job is queued: Create aggr1_controller_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE

```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate aggregate-name
```

Implementing the MetroCluster configuration

You must run the `metrocluster configure` command to start data protection in a MetroCluster configuration.

Before you begin

- There should be at least two non-root mirrored data aggregates on each cluster.

Additional data aggregates can be either mirrored or unmirrored.

Verify the aggregate types:

```
storage aggregate show
```



If you want to use a single mirrored data aggregate, then see [Configure MCC software in ONTAP](#) for instructions.

- The ha-config state of the controllers and chassis must be “mcc-2n”.

About this task

You can issue the `metrocluster configure` command once, on any of the nodes, to enable the MetroCluster configuration. You do not need to issue the command on each of the sites or nodes, and it does not matter which node or site you choose to issue the command on.

Steps

1. Configure the MetroCluster in the following format:

If your MetroCluster configuration has...	Then do this...
Multiple data aggregates	From any node’s prompt, configure MetroCluster: <code>metrocluster configure <i>node-name</i></code>

A single mirrored data aggregate

a. From any node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with "y" when you are prompted to continue into advanced mode and you see the advanced mode prompt (*>).

b. Configure the MetroCluster with the `-allow-with-one-aggregate true` parameter:

```
metrocluster configure -allow-with-one-aggregate true node-name
```

c. Return to the admin privilege level:

```
set -privilege admin
```



The best practice is to have multiple data aggregates. If the first DR group has only one aggregate and you want to add a DR group with one aggregate, you must move the metadata volume off the single data aggregate. For more information on this procedure, see [Moving a metadata volume in MetroCluster configurations](#).

The following command enables the MetroCluster configuration on all of the nodes in the DR group that contains "controller_A_1":

```
cluster_A::*> metrocluster configure -node-name controller_A_1  
  
[Job 121] Job succeeded: Configure is successful.
```

2. Verify the networking status on site A:

```
network port show
```

The following example shows the network port usage:

```
cluster_A::> network port show
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper

controller_A_1						
	e0a	Cluster	Cluster	up	9000	auto/1000
	e0b	Cluster	Cluster	up	9000	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
	e0g	Default	Default	up	1500	auto/1000

```
7 entries were displayed.
```

3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Verify the configuration from site A:

```
metrocluster show
```

```
cluster_A::> metrocluster show
```

Cluster	Entry Name	State

Local: cluster_A	Configuration state	configured
	Mode	normal
	AUSO Failure Domain	auso-on-cluster-
disaster		
Remote: cluster_B	Configuration state	configured
	Mode	normal
	AUSO Failure Domain	auso-on-cluster-
disaster		

b. Verify the configuration from site B:

```
metrocluster show
```

```

cluster_B::> metrocluster show
Cluster                               Entry Name                               State
-----                               -
Local: cluster_B                       Configuration state configured
Mode                                    normal
AUSO Failure Domain                    auso-on-cluster-
disaster
Remote: cluster_A                       Configuration state configured
Mode                                    normal
AUSO Failure Domain                    auso-on-cluster-
disaster

```

Configuring FC-to-SAS bridges for health monitoring

In systems running ONTAP versions prior to 9.8, if your configuration includes FC-to-SAS bridges, you must perform some special configuration steps to monitor the FC-to-SAS bridges in the MetroCluster configuration.

- Third-party SNMP monitoring tools are not supported for FibreBridge bridges.
- Beginning with ONTAP 9.8, FC-to-SAS bridges are monitored via in-band connections by default, and additional configuration is not required.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. From the ONTAP cluster prompt, add the bridge to health monitoring:
 - a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
ONTAP 9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
ONTAP 9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

- b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all of the data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```

controller_A_1::> storage bridge show

Bridge          Symbolic Name Is Monitored  Monitor Status
Vendor Model    Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01         true          ok            Atto
FibreBridge 7500N  20000010867038c0
ATTO_10.10.20.11  atto02         true          ok            Atto
FibreBridge 7500N  20000010867033c0
ATTO_10.10.20.12  atto03         true          ok            Atto
FibreBridge 7500N  20000010867030c0
ATTO_10.10.20.13  atto04         true          ok            Atto
FibreBridge 7500N  2000001086703b80

4 entries were displayed

controller_A_1::>

```

Checking the MetroCluster configuration

You can check that the components and relationships in the MetroCluster configuration are working correctly. You should do a check after initial configuration and after making any changes to the MetroCluster configuration. You should also do a check before a negotiated (planned) switchover or a switchback operation.

If the `metrocluster check run` command is issued twice within a short time on either or both clusters, a conflict can occur and the command might not collect all data. Subsequent `metrocluster check show` commands do not show the expected output.

1. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```

cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results. To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"

```

```
cluster_A::> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

2. Display more detailed results:

```
metrocluster check run
```

```
metrocluster check aggregate show
```

```
metrocluster check cluster show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```

The `metrocluster check show` commands show the results of the most recent `metrocluster check run` command. You should always run the `metrocluster check run` command prior to using the `metrocluster check show` commands so that the information displayed is current.

The following example shows the `metrocluster check aggregate show` command output for a healthy four-node MetroCluster configuration:

```
cluster_A::> metrocluster check aggregate show
```

```
Last Checked On: 8/5/2014 00:42:58
```

Node	Aggregate	Check
Result		
-----	-----	-----
controller_A_1	controller_A_1_aggr0	mirroring-status
ok		disk-pool-allocation
ok		

```

ok                                     ownership-state
                                     controller_A_1_aggr1
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_1_aggr2
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_2_aggr0
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_2_aggr1
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
ok                                     controller_A_2_aggr2
                                     mirroring-status
ok                                     disk-pool-allocation
ok                                     ownership-state
18 entries were displayed.

```

The following example shows the `metrocluster check cluster show` command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

Last Checked On: 9/13/2017 20:47:04

Cluster	Check	Result
mccint-fas9000-0102	negotiated-switchover-ready	not-applicable
	switchback-ready	not-applicable
	job-schedules	ok
	licenses	ok
	periodic-check-enabled	ok
mccint-fas9000-0304	negotiated-switchover-ready	not-applicable
	switchback-ready	not-applicable
	job-schedules	ok
	licenses	ok
	periodic-check-enabled	ok

10 entries were displayed.

Related information

[Disk and aggregate management](#)

[Network and LIF management](#)

Checking for MetroCluster configuration errors with Config Advisor

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

1. Go to the Config Advisor download page and download the tool.

[NetApp Downloads: Config Advisor](#)

2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Verifying switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

1. Use the procedures for negotiated switchover, healing, and switchback that are mentioned in the [Perform switchover, healing, and switchback](#).

Protecting configuration backup files

You can provide additional protection for the cluster configuration backup files by specifying a remote URL (either HTTP or FTP) where the configuration backup files will be uploaded in addition to the default locations in the local cluster.

1. Set the URL of the remote destination for the configuration backup files:

```
system configuration backup settings modify URL-of-destination
```

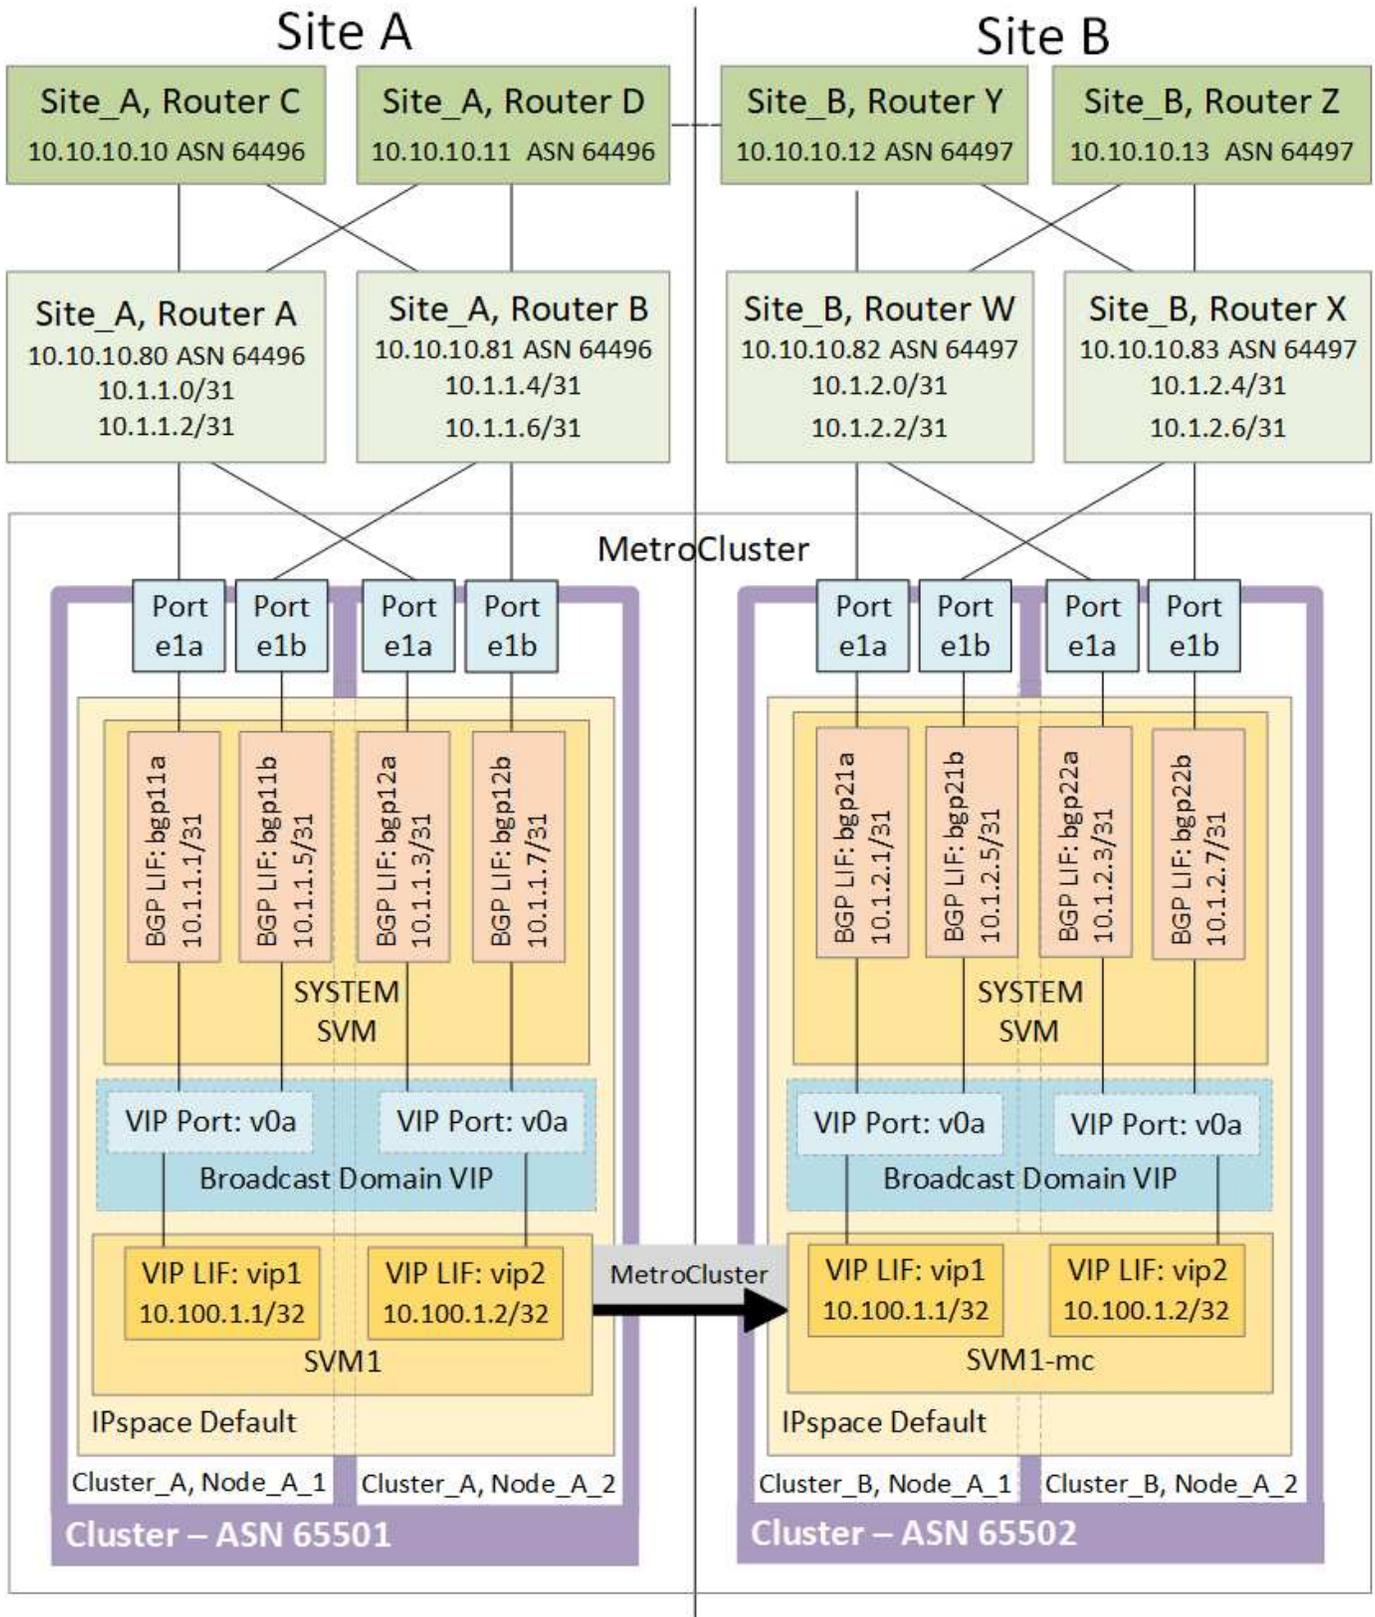
The [Cluster Management with the CLI](#) contains additional information under the section *Managing configuration backups*.

Considerations for using virtual IP and Border Gateway Protocol with a MetroCluster configuration

Beginning with ONTAP 9.5, ONTAP supports layer 3 connectivity using virtual IP (VIP) and Border Gateway Protocol (BGP). The combination VIP and BGP for redundancy in the front-end networking with the back-end MetroCluster redundancy provides a layer 3 disaster recovery solution.

Review the following guidelines and illustration when planning your layer 3 solution. For details on implementing VIP and BGP in ONTAP, refer to the following section:

[Configuring virtual IP \(VIP\) LIFs](#)



ONTAP limitations

ONTAP does not automatically verify that all nodes on both sites of the MetroCluster configuration are configured with BGP peering.

ONTAP does not perform route aggregation but announces all individual virtual LIF IPs as unique host routes

at all times.

ONTAP does not support true AnyCast — only a single node in the cluster presents a specific virtual LIF IP (but is accepted by all physical interfaces, regardless of whether they are BGP LIFs, provided the physical port is part of the correct IPspace). Different LIFs can migrate independently of each other to different hosting nodes.

Guidelines for using this Layer 3 solution with a MetroCluster configuration

You must configure your BGP and VIP correctly to provide the required redundancy.

Simpler deployment scenarios are preferred over more complex architectures (for example, a BGP peering router is reachable across an intermediate, non-BGP router). However, ONTAP does not enforce network design or topology restrictions.

VIP LIFs only cover the frontend/data network.

Depending on your version of ONTAP, you must configure BGP peering LIFs in the node SVM, not the system or data SVM. In ONTAP 9.8, the BGP LIFs are visible in the cluster (system) SVM and the node SVMs are no longer present.

Each data SVM requires the configuration of all potential first hop gateway addresses (typically, the BGP router peering IP address), so that the return data path is available if a LIF migration or MetroCluster failover occurs.

BGP LIFs are node specific, similar to intercluster LIFs — each node has a unique configuration, which does not need to be replicated to DR site nodes.

The existence of the v0a (v0b and so on.) continuously validates the connectivity, guaranteeing that a LIF migrate or failover succeeds (unlike L2, where a broken configuration is only visible after the outage).

A major architectural difference is that clients should no longer share the same IP subnet as the VIP of data SVMs. An L3 router with appropriate enterprise grade resiliency and redundancy features enabled (for example, VRRP/HSRP) should be on the path between storage and clients for the VIP to operate correctly.

The reliable update process of BGP allows for smoother LIF migrations because they are marginally faster and have a lower chance of interruption to some clients.

You can configure BGP to detect some classes of network or switch misbehaviors faster than LACP, if configured accordingly.

External BGP (EBGP) uses different AS numbers between ONTAP node(s) and peering routers and is the preferred deployment to ease route aggregation and redistribution on the routers. Internal BGP (IBGP) and the use of route reflectors is not impossible but outside the scope of a straightforward VIP setup.

After deployment, you must check that the data SVM is accessible when the associated virtual LIF is migrated between all nodes on each site (including MetroCluster switchover) to verify the correct configuration of the static routes to the same data SVM.

VIP works for most IP-based protocols (NFS, SMB, iSCSI).

Testing the MetroCluster configuration

You can test failure scenarios to confirm the correct operation of the MetroCluster configuration.

Verifying negotiated switchover

You can test a negotiated (planned) switchover operation to confirm uninterrupted data availability.

This test validates that data availability is not affected (except for SMB and Fibre Channel protocols) by switching the cluster over to the second data center.

This test should take about 30 minutes.

This procedure has the following expected results:

- The `metrocluster switchover` command will present a warning prompt.

If you respond **yes** to the prompt, the site the command is issued from will switch over the partner site.

For MetroCluster IP configurations:

- For ONTAP 9.4 and earlier:
 - Mirrored aggregates will become degraded after the negotiated switchover.
- For ONTAP 9.5 and later:
 - Mirrored aggregates will remain in normal state if the remote storage is accessible.
 - Mirrored aggregates will become degraded after the negotiated switchover if access to the remote storage is lost.
- For ONTAP 9.8 and later:
 - Unmirrored aggregates that are located at the disaster site will become unavailable if access to the remote storage is lost. This might lead to a controller outage.

Steps

1. Confirm that all nodes are in the configured state and normal mode:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show

Cluster                               Configuration State  Mode
-----
Local: cluster_A                       configured           normal
Remote: cluster_B                       configured           normal
```

2. Begin the switchover operation:

```
metrocluster switchover
```

```
cluster_A::> metrocluster switchover
Warning: negotiated switchover is about to start. It will stop all the
data Vservers on cluster "cluster_B" and
automatically re-start them on cluster "cluster_A". It will finally
gracefully shutdown cluster "cluster_B".
```

3. Confirm that the local cluster is in the configured state and switchover mode:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show

Cluster                Configuration State  Mode
-----
Local: cluster_A       configured           switchover
Remote: cluster_B      not-reachable       -
                       configured           normal
```

4. Confirm that the switchover operation was successful:

```
metrocluster operation show
```

```
cluster_A::> metrocluster operation show
Operation: switchover
State: successful
Start Time: 2/6/2016 13:28:50
End Time: 2/6/2016 13:29:41
Errors: -
```

5. Use the `vserver show` and `network interface show` commands to verify that DR SVMs and LIFs have come online.

Verifying healing and manual switchback

You can test the healing and manual switchback operations to verify that data availability is not affected (except for SMB and Solaris FC configurations) by switching back the cluster to the original data center after a negotiated switchover.

This test should take about 30 minutes.

The expected result of this procedure is that services should be switched back to their home nodes.

Steps

1. Verify that healing is completed:

```
metrocluster node show
```

The following example shows the successful completion of the command:

```
cluster_A::> metrocluster node show
DR
Group Cluster Node          Configuration  DR
                    State          Mirroring Mode
-----
-----
1      cluster_A
           node_A_1      configured    enabled    heal roots
completed
           cluster_B
           node_B_2      unreachable  -          switched over
42 entries were displayed.
```

2. Verify that all aggregates are mirrored:

```
storage aggregate show
```

The following example shows that all aggregates have a RAID Status of mirrored:

```

cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate Size      Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
data_cluster
      4.19TB      4.13TB   2% online    8 node_A_1  raid_dp,
mirrored,
normal

root_cluster
      715.5GB     212.7GB  70% online    1 node_A_1  raid4,
mirrored,
normal

cluster_B Switched Over Aggregates:
Aggregate Size      Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
data_cluster_B
      4.19TB      4.11TB   2% online    5 node_A_1  raid_dp,
mirrored,
normal

root_cluster_B     -          -      - unknown    - node_A_1  -

```

3. Boot nodes from the disaster site.
4. Check the status of switchback recovery:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR
Group Cluster Node      Configuration  DR
State          Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured    enabled    heal roots
completed
      cluster_B
      node_B_2      configured    enabled    waiting for
switchback                                     recovery

2 entries were displayed.

```

5. Perform the switchback:

```
metrocluster switchback
```

```
cluster_A::> metrocluster switchback
[Job 938] Job succeeded: Switchback is successful. Verify switchback
```

6. Confirm status of the nodes:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
DR
Group Cluster Node Configuration State DR Mirroring Mode
-----
1 cluster_A
node_A_1 configured enabled normal
cluster_B
node_B_2 configured enabled normal
2 entries were displayed.
```

7. Confirm the status:

```
metrocluster operation show
```

The output should show a successful state.

```
cluster_A::> metrocluster operation show
Operation: switchback
State: successful
Start Time: 2/6/2016 13:54:25
End Time: 2/6/2016 13:56:15
Errors: -
```

Loss of a single FC-to-SAS bridge

You can test the failure of a single FC-to-SAS bridge to make sure there is no single point of failure.

This test should take about 15 minutes.

This procedure has the following expected results:

- Errors should be generated as the bridge is switched off.

- No failover or loss of service should occur.
- Only one path from the controller module to the drives behind the bridge is available.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. Turn off the power supplies of the bridge.
2. Confirm that the bridge monitoring indicates an error:

```
storage bridge show
```

```
cluster_A::> storage bridge show

Monitor
Bridge      Symbolic Name Vendor  Model      Bridge WWN      Monitored
Status
-----
-----
ATTO_10.65.57.145
      bridge_A_1   Atto    FibreBridge 6500N
                                      200000108662d46c true
error
```

3. Confirm that drives behind the bridge are available with a single path:

```
storage disk error show
```

```

cluster_A::> storage disk error show
Disk              Error Type          Error Text
-----
-----
1.0.0             onedomain           1.0.0 (5000cca057729118): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.
1.0.1             onedomain           1.0.1 (5000cca057727364): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.
1.0.2             onedomain           1.0.2 (5000cca05772e9d4): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.
...
1.0.23            onedomain           1.0.23 (5000cca05772e9d4): All paths
to this array LUN are connected to the same fault domain. This is a
single point of failure.

```

Verifying operation after power line disruption

You can test the MetroCluster configuration's response to the failure of a PDU.

The best practice is for each power supply unit (PSU) in a component to be connected to a separate power supply. If both PSUs are connected to the same power distribution unit (PDU) and an electrical disruption occurs, the site could down and a complete shelf might become unavailable. Failure of one power line is tested to confirm that there is no cabling mismatch that could cause a service disruption.

This test should take about 15 minutes.

This test requires turning off power to all left-hand PDUs and then all right-hand PDUs on all of the racks containing the MetroCluster components.

This procedure has the following expected results:

- Errors should be generated as the PDUs are disconnected.
- No failover or loss of service should occur.

Steps

1. Turn off the power of the PDUs on the left-hand side of the rack containing the MetroCluster components.
2. Monitor the result on the console by using the `system environment sensors show -state fault` and `storage shelf show -errors` commands.

```

cluster_A::> system environment sensors show -state fault

Node Sensor                State Value/Units Crit-Low Warn-Low Warn-Hi
Crit-Hi
-----
node_A_1
    PSU1                    fault
                               PSU_OFF
    PSU1 Pwr In OK          fault
                               FAULT
node_A_2
    PSU1                    fault
                               PSU_OFF
    PSU1 Pwr In OK          fault
                               FAULT

4 entries were displayed.

cluster_A::> storage shelf show -errors
    Shelf Name: 1.1
    Shelf UID: 50:0a:09:80:03:6c:44:d5
    Serial Number: SHFHU1443000059

Error Type                Description
-----
Power                      Critical condition is detected in storage shelf
power supply unit "1". The unit might fail.Reconnect PSU1

```

3. Turn the power back on to the left-hand PDUs.
4. Make sure that ONTAP clears the error condition.
5. Repeat the previous steps with the right-hand PDUs.

Verifying operation after loss of a single storage shelf

You can test the failure of a single storage shelf to verify that there is no single point of failure.

This procedure has the following expected results:

- An error message should be reported by the monitoring software.
- No failover or loss of service should occur.
- Mirror resynchronization starts automatically after the hardware failure is restored.

Steps

1. Check the storage failover status:

```
storage failover show
```

```
cluster_A::> storage failover show
```

Node	Partner	Possible	State	Description
node_A_1	node_A_2	true	Connected to	node_A_2
node_A_2	node_A_1	true	Connected to	node_A_1

2 entries were displayed.

2. Check the aggregate status:

```
storage aggregate show
```

```
cluster_A::> storage aggregate show
```

```
cluster Aggregates:
```

```
Aggregate      Size Available Used% State   #Vols  Nodes      RAID  
Status
```

```
-----  
-----
```

```
node_A_1data01_mirrored  
      4.15TB    3.40TB    18% online    3 node_A_1
```

```
raid_dp,
```

```
mirrored,
```

```
normal
```

```
node_A_1root  
      707.7GB   34.29GB   95% online    1 node_A_1
```

```
raid_dp,
```

```
mirrored,
```

```
normal
```

```
node_A_2_data01_mirrored  
      4.15TB    4.12TB    1% online    2 node_A_2
```

```
raid_dp,
```

```
mirrored,
```

```
normal
```

```
node_A_2_data02_unmirrored  
      2.18TB    2.18TB    0% online    1 node_A_2
```

```
raid_dp,
```

```
normal
```

```
node_A_2_root  
      707.7GB   34.27GB   95% online    1 node_A_2
```

```
raid_dp,
```

```
mirrored,
```

```
normal
```

3. Verify that all data SVMs and data volumes are online and serving data:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0,!MDV*
```

```
cluster_A::> vserver show -type data
```

```
cluster_A::> vserver show -type data
```

Vserver	Type	Subtype	Admin State	Operational State	Root Volume
Aggregate					
SVM1	data	sync-source		running	SVM1_root
node_A_1_data01_mirrored					
SVM2	data	sync-source		running	SVM2_root
node_A_2_data01_mirrored					

```
cluster_A::> network interface show -fields is-home false
```

```
There are no entries matching your query.
```

```
cluster_A::> volume show !vol0,!MDV*
```

Vserver	Volume	Aggregate	State	Type	Size
Available	Used%				
SVM1					
	SVM1_root	node_A_1data01_mirrored	online	RW	10GB
9.50GB	5%				
SVM1					
	SVM1_data_vol	node_A_1data01_mirrored	online	RW	10GB
9.49GB	5%				
SVM2					
	SVM2_root	node_A_2_data01_mirrored	online	RW	10GB
9.49GB	5%				
SVM2					
	SVM2_data_vol	node_A_2_data02_unmirrored	online	RW	1GB
972.6MB	5%				

4. Identify a shelf in Pool 1 for node node_A_2 to power off to simulate a sudden hardware failure:

```
storage aggregate show -r -node node-name !*root
```

The shelf you select must contain drives that are part of a mirrored data aggregate.

In the following example, shelf ID 31 is selected to fail.

```
cluster_A::> storage aggregate show -r -node node_A_2 !*root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirrored) (block
checksums)
Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block
checksums)
```

Physical	Position	Disk	Pool	Type	RPM	Usable
Size	Status					Size
828.0GB (normal)	dparity	2.30.3	0	BSAS	7200	827.7GB
828.0GB (normal)	parity	2.30.4	0	BSAS	7200	827.7GB
828.0GB (normal)	data	2.30.6	0	BSAS	7200	827.7GB
828.0GB (normal)	data	2.30.8	0	BSAS	7200	827.7GB
828.0GB (normal)	data	2.30.5	0	BSAS	7200	827.7GB

```

Plex: /node_A_2_data01_mirrored/plex4 (online, normal, active, pool1)
RAID Group /node_A_2_data01_mirrored/plex4/rg0 (normal, block
checksums)
```

Physical	Position	Disk	Pool	Type	RPM	Usable
Size	Status					Size
828.0GB (normal)	dparity	1.31.7	1	BSAS	7200	827.7GB
828.0GB (normal)	parity	1.31.6	1	BSAS	7200	827.7GB
828.0GB (normal)	data	1.31.3	1	BSAS	7200	827.7GB

```

    data      1.31.4                1   BSAS      7200  827.7GB
828.0GB (normal)
    data      1.31.5                1   BSAS      7200  827.7GB
828.0GB (normal)

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block
checksums)
Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active,
pool0)
RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block
checksums)

                                                    Usable
Physical
  Position Disk                               Pool Type    RPM    Size
Size Status
-----
-----
    dparity  2.30.12                0   BSAS      7200  827.7GB
828.0GB (normal)
    parity   2.30.22                0   BSAS      7200  827.7GB
828.0GB (normal)
    data     2.30.21                0   BSAS      7200  827.7GB
828.0GB (normal)
    data     2.30.20                0   BSAS      7200  827.7GB
828.0GB (normal)
    data     2.30.14                0   BSAS      7200  827.7GB
828.0GB (normal)
15 entries were displayed.

```

5. Physically power off the shelf that you selected.

6. Check the aggregate status again:

```
storage aggregate
```

```
storage aggregate show -r -node node_A_2 !*root
```

The aggregate with drives on the powered-off shelf should have a “degraded” RAID status, and drives on the affected plex should have a “failed” status, as shown in the following example:

```

cluster_A::> storage aggregate show
Aggregate      Size Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
node_A_1data01_mirrored
                4.15TB    3.40TB   18% online    3 node_A_1

```

```

raid_dp,

mirrored,

normal
node_A_1root
          707.7GB   34.29GB   95% online      1 node_A_1

```

```

raid_dp,

mirrored,

normal
node_A_2_data01_mirrored
          4.15TB    4.12TB    1% online      2 node_A_2

```

```

raid_dp,

mirror

degraded
node_A_2_data02_unmirrored
          2.18TB    2.18TB    0% online      1 node_A_2

```

```

raid_dp,

normal
node_A_2_root
          707.7GB   34.27GB   95% online      1 node_A_2

```

```

raid_dp,

mirror

degraded
cluster_A::> storage aggregate show -r -node node_A_2 !*root
Owner Node: node_A_2
Aggregate: node_A_2_data01_mirrored (online, raid_dp, mirror degraded)
(block checksums)
Plex: /node_A_2_data01_mirrored/plex0 (online, normal, active, pool0)
RAID Group /node_A_2_data01_mirrored/plex0/rg0 (normal, block
checksums)

```

				Usable	
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	2.30.3	0	BSAS	7200	827.7GB
828.0GB (normal)					

```

    parity    2.30.4                0    BSAS    7200    827.7GB
828.0GB (normal)
    data      2.30.6                0    BSAS    7200    827.7GB
828.0GB (normal)
    data      2.30.8                0    BSAS    7200    827.7GB
828.0GB (normal)
    data      2.30.5                0    BSAS    7200    827.7GB
828.0GB (normal)

```

Plex: /node_A_2_data01_mirrored/plex4 (offline, failed, inactive, pool1)

RAID Group /node_A_2_data01_mirrored/plex4/rg0 (partial, none checksums)

					Usable
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	FAILED	-	-	-	827.7GB
- (failed)					
parity	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					
data	FAILED	-	-	-	827.7GB
- (failed)					

Aggregate: node_A_2_data02_unmirrored (online, raid_dp) (block checksums)

Plex: /node_A_2_data02_unmirrored/plex0 (online, normal, active, pool0)

RAID Group /node_A_2_data02_unmirrored/plex0/rg0 (normal, block checksums)

					Usable
Physical					
Position	Disk	Pool	Type	RPM	Size
Size	Status				

dparity	2.30.12	0	BSAS	7200	827.7GB
828.0GB (normal)					
parity	2.30.22	0	BSAS	7200	827.7GB
828.0GB (normal)					

```
data      2.30.21      0  BSAS  7200  827.7GB
828.0GB (normal)
data      2.30.20      0  BSAS  7200  827.7GB
828.0GB (normal)
data      2.30.14      0  BSAS  7200  827.7GB
828.0GB (normal)
```

15 entries were displayed.

7. Verify that the data is being served and that all volumes are still online:

```
vserver show -type data
```

```
network interface show -fields is-home false
```

```
volume show !vol0,!MDV*
```

```

cluster_A::> vservers show -type data

cluster_A::> vservers show -type data
Admin      Operational Root
Vserver    Type      Subtype    State      State      Volume
Aggregate
-----
-----
SVM1       data      sync-source  running    SVM1_root
node_A_1_data01_mirrored
SVM2       data      sync-source  running    SVM2_root
node_A_1_data01_mirrored

cluster_A::> network interface show -fields is-home false
There are no entries matching your query.

cluster_A::> volume show !vol0,!MDV*
Vserver    Volume      Aggregate    State      Type      Size
Available Used%
-----
-----
SVM1
          SVM1_root
                node_A_1data01_mirrored
                        online      RW      10GB
9.50GB    5%
SVM1
          SVM1_data_vol
                node_A_1data01_mirrored
                        online      RW      10GB
9.49GB    5%
SVM2
          SVM2_root
                node_A_1data01_mirrored
                        online      RW      10GB
9.49GB    5%
SVM2
          SVM2_data_vol
                node_A_2_data02_unmirrored
                        online      RW      1GB
972.6MB   5%

```

8. Physically power on the shelf.

Resynchronization starts automatically.

9. Verify that resynchronization has started:

```
storage aggregate show
```

The affected aggregate should have a “resyncing” RAID status, as shown in the following example:

```
cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
node_A_1_data01_mirrored
      4.15TB      3.40TB      18% online      3 node_A_1
raid_dp,
mirrored,
normal
node_A_1_root
      707.7GB      34.29GB      95% online      1 node_A_1
raid_dp,
mirrored,
normal
node_A_2_data01_mirrored
      4.15TB      4.12TB       1% online      2 node_A_2
raid_dp,
resyncing
node_A_2_data02_unmirrored
      2.18TB      2.18TB       0% online      1 node_A_2
raid_dp,
normal
node_A_2_root
      707.7GB      34.27GB      95% online      1 node_A_2
raid_dp,
resyncing
```

10. Monitor the aggregate to confirm that resynchronization is complete:

```
storage aggregate show
```

The affected aggregate should have a “normal” RAID status, as shown in the following example:

```

cluster_A::> storage aggregate show
cluster Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
node_A_1data01_mirrored
          4.15TB      3.40TB   18% online    3 node_A_1
raid_dp,
mirrored,
normal
node_A_1root
          707.7GB    34.29GB   95% online    1 node_A_1
raid_dp,
mirrored,
normal
node_A_2_data01_mirrored
          4.15TB      4.12TB    1% online    2 node_A_2
raid_dp,
normal
node_A_2_data02_unmirrored
          2.18TB      2.18TB    0% online    1 node_A_2
raid_dp,
normal
node_A_2_root
          707.7GB    34.27GB   95% online    1 node_A_2
raid_dp,
resyncing

```

Remove MetroCluster configurations

If you need to remove the MetroCluster configuration, contact technical support.

Contact NetApp technical support and reference the appropriate guide for your configuration from [How to remove nodes from a MetroCluster configuration - Resolution Guide](#).



You cannot reverse the MetroCluster unconfiguration. This process should only be done with the assistance of technical support. After removing the MetroCluster configuration, all disk connectivity and interconnects should be adjusted to be in a supported state.

How to use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

Use the Active IQ Unified Manager and ONTAP System Manager for further configuration and monitoring

The Active IQ Unified Manager and ONTAP System Manager can be used for GUI management of the clusters and monitoring the configuration.

Each node has ONTAP System Manager pre-installed. To load System Manager, enter the cluster management LIF address as the URL in a web browser that has connectivity to the node.

You can also use Active IQ Unified Manager to monitor the MetroCluster configuration.

Related information

[Active IQ Unified Manager Documentation](#)

Synchronize the system time using NTP

Each cluster needs its own Network Time Protocol (NTP) server to synchronize the time between the nodes and their clients.

About this task

- You cannot modify the time zone settings for a failed node or the partner node after takeover occurs.
- Each cluster in the stretch MetroCluster configuration should have its own separate NTP server or servers used by the nodes at that MetroCluster site.
- If you are using the MetroCluster Tiebreaker software, it should also have its own separate NTP server.

Depending on your ONTAP version, you can configure the NTP from the **Cluster** or **Insights** tab in the System Manager UI.

Cluster

In System Manager, you can configure the NTP from the **Cluster** tab using two different options, depending on your ONTAP version:

ONTAP 9.8 or later:

Use the following steps to synchronize the NTP from the **Cluster** tab in ONTAP 9.8 or later.

Steps

1. Go to **Cluster > Overview**
2. Then select the  **More** option and select **Edit**.
3. In the **Edit Cluster Details** window, select the **+Add** option below NTP Servers.
4. Add the name, location, and specify the IP address of the time server.
5. Then, select **Save**.
6. Repeat the steps for any additional time servers.

ONTAP 9.11.1 or later:

Use the following steps to synchronize the NTP from the **Insights** window in the **Cluster** tab in ONTAP 9.11.1 or later.

Steps

1. Go to **Cluster > Overview**
2. Scroll down to the **Insights** window on the page, locate **Too few NTP servers are configured**, and then select **Fix It**.
3. Specify the IP address of the time server, and then select **Save**.
4. Repeat the previous step for any additional time servers.

Insights

In ONTAP 9.11.1 or later, you can also configure the NTP by using the **Insights** tab in System Manager:

Steps

1. Go to the **Insights** tab in the System Manager UI.
2. Scroll down to **Too few NTP servers are configured** and select **Fix It**.
3. Specify the IP address of the time server, and then select **Save**.
4. Repeat the previous step for any additional time servers.

Considerations when using ONTAP in a MetroCluster configuration

When using ONTAP in a MetroCluster configuration, you should be aware of certain considerations for licensing, peering to clusters outside the MetroCluster configuration, performing volume operations, NVFAIL operations, and other ONTAP operations.

Licensing considerations

- Both sites should be licensed for the same site-licensed features.
- All nodes should be licensed for the same node-locked features.

SnapMirror consideration

- SnapMirror SVM disaster recovery is only supported on MetroCluster configurations running versions of ONTAP 9.5 or later.

FlexCache support in a MetroCluster configuration

Beginning with ONTAP 9.7, FlexCache volumes are supported on MetroCluster configurations. You should be aware of requirements for manual repeer after switchover or switchback operations.

SVM repeer after switchover when FlexCache origin and cache are within the same MetroCluster site

After a negotiated or unplanned switchover, any SVM FlexCache peering relationship within the cluster must be manually configured.

For example, SVMs vs1 (cache) and vs2 (origin) are on site_A. These SVMs are peered.

After switchover, SVMs vs1-mc and vs2-mc are activated at the partner site (site_B). They must be manually repeer for FlexCache to work using the `vserver peer repeer` command.

SVM repeer after switchover or switchback when a FlexCache destination is on a third cluster and in disconnected mode

For FlexCache relationships to a cluster outside of the MetroCluster configuration, the peering must always be manually reconfigured after a switchover when the clusters involved are in a disconnected mode during switchover.

For example:

- One end of the FlexCache (cache_1 on vs1) resides on MetroCluster site_A has one end of the FlexCache
- The other end of the FlexCache (origin_1 on vs2) resides on site_C (not in the MetroCluster configuration)

When switchover is triggered, and if site_A and site_C are not connected, you must manually repeer the SVMs on site_B (the switchover cluster) and site_C using the `vserver peer repeer` command after the switchover.

When switchback is performed, you must again repeer the SVMs on site_A (the original cluster) and site_C.

FabricPool support in MetroCluster configurations

Beginning with ONTAP 9.7, MetroCluster configurations support FabricPool storage tiers.

For general information on using FabricPools, see the [Disks and aggregates management](#).

Considerations when using FabricPools

- The clusters must have FabricPool licenses with matching capacity limits.

- The clusters must have IPspaces with matching names.

This can be the default IPspace, or an IP space an administrator has created. This IPspace will be used for FabricPool object store configuration setups.

- For the selected IPspace, each cluster must have an intercluster LIF defined that can reach the external object store.

Configuring an aggregate for use in a mirrored FabricPool



Before you configure the aggregate you must set up object stores as described in "Setting up object stores for FabricPool in a MetroCluster configuration" in the [Disks and aggregates management](#).

To configure an aggregate for use in a FabricPool:

1. Create the aggregate or select an existing aggregate.
2. Mirror the aggregate as a typical mirrored aggregate within the MetroCluster configuration.
3. Create the FabricPool mirror with the aggregate, as described in the [Disks and aggregates management](#):
 - a. Attach a primary object store.

This object store is physically closer to the cluster.

- b. Add a mirror object store.

This object store is physically further away from the cluster than the primary object store.

FlexGroup support in MetroCluster configurations

Beginning with ONTAP 9.6 MetroCluster configurations support FlexGroup volumes.

Job schedules in a MetroCluster configuration

In ONTAP 9.3 and later, user-created job schedules are automatically replicated between clusters in a MetroCluster configuration. If you create, modify, or delete a job schedule on a cluster, the same schedule is automatically created on the partner cluster, using Configuration Replication Service (CRS).



System-created schedules are not replicated and you must manually perform the same operation on the partner cluster so that job schedules on both clusters are identical.

Cluster peering from the MetroCluster site to a third cluster

Because the peering configuration is not replicated, if you peer one of the clusters in the MetroCluster configuration to a third cluster outside of that configuration, you must also configure the peering on the partner MetroCluster cluster. This is so that peering can be maintained if a switchover occurs.

The non-MetroCluster cluster must be running ONTAP 8.3 or later. If not, peering is lost if a switchover occurs even if the peering has been configured on both MetroCluster partners.

LDAP client configuration replication in a MetroCluster configuration

An LDAP client configuration created on a storage virtual machine (SVM) on a local cluster is replicated to its partner data SVM on the remote cluster. For example, if the LDAP client configuration is created on the admin SVM on the local cluster, then it is replicated to all the admin data SVMs on the remote cluster. This MetroCluster feature is intentional so that the LDAP client configuration is active on all the partner SVMs on the remote cluster.

Networking and LIF creation guidelines for MetroCluster configurations

You should be aware of how LIFs are created and replicated in a MetroCluster configuration. You must also know about the requirement for consistency so that you can make proper decisions when configuring your network.

Related information

[ONTAP concepts](#)

IPspace object replication and subnet configuration requirements

You should be aware of the requirements for replicating IPspace objects to the partner cluster and for configuring subnets and IPv6 in a MetroCluster configuration.

IPspace replication

You must consider the following guidelines while replicating IPspace objects to the partner cluster:

- The IPspace names of the two sites must match.
- IPspace objects must be manually replicated to the partner cluster.

Any storage virtual machines (SVMs) that are created and assigned to an IPspace before the IPspace is replicated will not be replicated to the partner cluster.

Subnet configuration

You must consider the following guidelines while configuring subnets in a MetroCluster configuration:

- Both clusters of the MetroCluster configuration must have a subnet in the same IPspace with the same subnet name, subnet, broadcast domain, and gateway.
- The IP ranges of the two clusters must be different.

In the following example, the IP ranges are different:

```
cluster_A::> network subnet show
```

```
IPspace: Default
```

Subnet		Broadcast		Avail/	
Name	Subnet	Domain	Gateway	Total	Ranges
-----	-----	-----	-----	-----	

subnet1	192.168.2.0/24	Default	192.168.2.1	10/10	
	192.168.2.11-192.168.2.20				

```
cluster_B::> network subnet show
```

```
IPspace: Default
```

Subnet		Broadcast		Avail/	
Name	Subnet	Domain	Gateway	Total	Ranges
-----	-----	-----	-----	-----	

subnet1	192.168.2.0/24	Default	192.168.2.1	10/10	
	192.168.2.21-192.168.2.30				

IPv6 configuration

If IPv6 is configured on one site, IPv6 must be configured on the other site as well.

Requirements for LIF creation in a MetroCluster configuration

You should be aware of the requirements for creating LIFs when configuring your network in a MetroCluster configuration.

You must consider the following guidelines when creating LIFs:

- Fibre Channel: You must use stretched VSAN or stretched fabrics.
- IP/iSCSI: You must use layer 2 stretched network.
- ARP broadcasts: You must enable ARP broadcasts between the two clusters.
- Duplicate LIFs: You must not create multiple LIFs with the same IP address (duplicate LIFs) in an IPspace.
- NFS and SAN configurations: You must use different storage virtual machines (SVMs) for both the unmirrored and mirrored aggregates.
- You should create a subnet object before you create a LIF. A subnet object enables ONTAP to determine failover targets on the destination cluster because it has an associated broadcast domain.

Verify LIF creation

You can confirm the successful creation of a LIF in a MetroCluster configuration by running the `metrocluster check lif show` command. If you encounter any issues while creating the LIF, you can use the `metrocluster check lif repair-placement` command to fix the issues.

LIF replication and placement requirements and issues

You should be aware of the LIF replication requirements in a MetroCluster configuration. You should also know how a replicated LIF is placed on a partner cluster, and you should be aware of the issues that occur when LIF replication or LIF placement fails.

Replication of LIFs to the partner cluster

When you create a LIF on a cluster in a MetroCluster configuration, the LIF is replicated on the partner cluster. LIFs are not placed on a one-to-one name basis. For availability of LIFs after a switchover operation, the LIF placement process verifies that the ports are able to host the LIF based on reachability and port attribute checks.

The system must meet the following conditions to place the replicated LIFs on the partner cluster:

Condition	LIF type: FC	LIF type: IP/iSCSI
Node identification	<p>ONTAP attempts to place the replicated LIF on the disaster recovery (DR) partner of the node on which it was created.</p> <p>If the DR partner is unavailable, the DR auxiliary partner is used for placement.</p>	<p>ONTAP attempts to place the replicated LIF on the DR partner of the node on which it was created.</p> <p>If the DR partner is unavailable, the DR auxiliary partner is used for placement.</p>
Port identification	<p>ONTAP identifies the connected FC target ports on the DR cluster.</p>	<p>The ports on the DR cluster that are in the same IPspace as the source LIF are selected for a reachability check.</p> <p>If there are no ports in the DR cluster in the same IPspace, the LIF cannot be placed.</p> <p>All of the ports in the DR cluster that are already hosting a LIF in the same IPspace and subnet are automatically marked as reachable; and can be used for placement. These ports are not included in the reachability check.</p>

Reachability check	<p>Reachability is determined by checking for the connectivity of the source fabric WWN on the ports in the DR cluster.</p> <p>If the same fabric is not present at the DR site, the LIF is placed on a random port on the DR partner.</p>	<p>Reachability is determined by the response to an Address Resolution Protocol (ARP) broadcast from each previously identified port on the DR cluster to the source IP address of the LIF to be placed.</p> <p>For reachability checks to succeed, ARP broadcasts must be allowed between the two clusters.</p> <p>Each port that receives a response from the source LIF will be marked as possible for placement.</p>
Port selection	<p>ONTAP categorizes the ports based on attributes such as adapter type and speed, and then selects the ports with matching attributes.</p> <p>If no ports with matching attributes are found, the LIF is placed on a random connected port on the DR partner.</p>	<p>From the ports that are marked as reachable during the reachability check, ONTAP prefers ports that are in the broadcast domain that is associated with the subnet of the LIF.</p> <p>If there are no network ports available on the DR cluster that are in the broadcast domain that is associated with the subnet of the LIF, then ONTAP selects ports that have reachability to the source LIF.</p> <p>If there are no ports with reachability to the source LIF, a port is selected from the broadcast domain that is associated with the subnet of the source LIF, and if no such broadcast domain exists, a random port is selected.</p> <p>ONTAP categorizes the ports based on attributes such as adapter type, interface type, and speed, and then selects the ports with matching attributes.</p>
LIF placement	From the reachable ports, ONTAP selects the least loaded port for placement.	From the selected ports, ONTAP selects the least loaded port for placement.

Placement of replicated LIFs when the DR partner node is down

When an iSCSI or FC LIF is created on a node whose DR partner has been taken over, the replicated LIF is placed on the DR auxiliary partner node. After a subsequent giveback operation, the LIFs are not automatically moved to the DR partner. This can lead to LIFs being concentrated on a single node in the partner cluster. During a MetroCluster switchover operation, subsequent attempts to map LUNs belonging to the storage

virtual machine (SVM) fail.

You should run the `metrocluster check lif show` command after a takeover operation or giveback operation to verify that the LIF placement is correct. If errors exist, you can run the `metrocluster check lif repair-placement` command to resolve the issues.

LIF placement errors

LIF placement errors that are displayed by the `metrocluster check lif show` command are retained after a switchover operation. If the `network interface modify`, `network interface rename`, or `network interface delete` command is issued for a LIF with a placement error, the error is removed and does not appear in the output of the `metrocluster check lif show` command.

LIF replication failure

You can also check whether LIF replication was successful by using the `metrocluster check lif show` command. An EMS message is displayed if LIF replication fails.

You can correct a replication failure by running the `metrocluster check lif repair-placement` command for any LIF that fails to find a correct port. You should resolve any LIF replication failures as soon as possible to verify the availability of LIF during a MetroCluster switchover operation.



Even if the source SVM is down, LIF placement might proceed normally if there is a LIF belonging to a different SVM in a port with the same IPspace and network in the destination SVM.

Volume creation on a root aggregate

The system does not allow the creation of new volumes on the root aggregate (an aggregate with an HA policy of CFO) of a node in a MetroCluster configuration.

Because of this restriction, root aggregates cannot be added to an SVM using the `vserver add-aggregates` command.

SVM disaster recovery in a MetroCluster configuration

Beginning with ONTAP 9.5, active storage virtual machines (SVMs) in a MetroCluster configuration can be used as sources with the SnapMirror SVM disaster recovery feature. The destination SVM must be on the third cluster outside of the MetroCluster configuration.

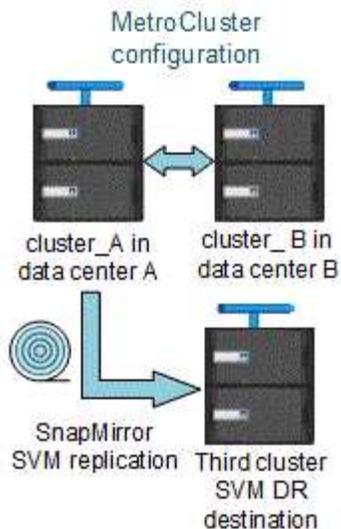
You should be aware of the following requirements and limitations of using SVMs with SnapMirror disaster recovery:

- Only an active SVM within a MetroCluster configuration can be the source of an SVM disaster recovery relationship.

A source can be a sync-source SVM before switchover or a sync-destination SVM after switchover.

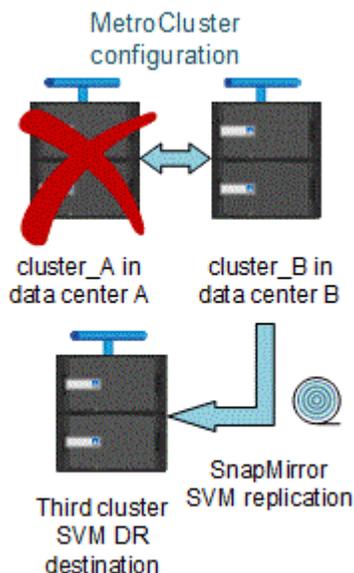
- When a MetroCluster configuration is in a steady state, the MetroCluster sync-destination SVM cannot be the source of an SVM disaster recovery relationship, since the volumes are not online.

The following image shows the SVM disaster recovery behavior in a steady state:



- When the sync-source SVM is the source of an SVM DR relationship, the source SVM DR relationship information is replicated to the MetroCluster partner.

This enables the SVM DR updates to continue after a switchover as shown in the following image:



- During the switchover and switchback processes, replication to the SVM DR destination might fail.

However, after the switchover or switchback process completes, the next SVM DR scheduled updates will succeed.

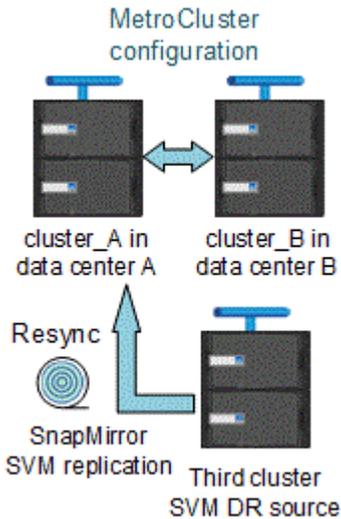
See the section “Replicating the SVM configuration” in the [Data Protection with the CLI](#) for details on configuring an SVM DR relationship.

SVM resynchronization at a disaster recovery site

During resynchronization, the storage virtual machines (SVMs) disaster recovery (DR) source on the MetroCluster configuration is restored from the destination SVM on the non-MetroCluster site.

During resynchronization, the source SVM (cluster_A) temporarily acts as a destination SVM as shown in the

following image:



If an unplanned switchover occurs during resynchronization

Unplanned switchovers that occur during the resynchronization will halt the resynchronization transfer. If an unplanned switchover occurs, the following conditions are true:

- The destination SVM on the MetroCluster site (which was a source SVM prior to resynchronization) remains as a destination SVM. The SVM at the partner cluster will continue to retain its subtype and remain inactive.
- The SnapMirror relationship must be re-created manually with the sync-destination SVM as the destination.
- The SnapMirror relationship does not appear in the SnapMirror show output after a switchover at the survivor site unless a SnapMirror create operation is executed.

Performing switchback after an unplanned switchover during resynchronization

To successfully perform the switchback process, the resynchronization relationship must be broken and deleted. Switchback is not permitted if there are any SnapMirror DR destination SVMs in the MetroCluster configuration or if the cluster has an SVM of subtype "dp-destination".

Output of the storage disk show and storage shelf show commands in a two-node stretch MetroCluster configuration

In a two-node stretch MetroCluster configuration, the `is-local-attach` field of the `storage disk show` and `storage shelf show` commands shows all of the disks and storage shelves as local, regardless of the node to which they are attached.

Output for the storage aggregate plex show command is indeterminate after a MetroCluster switchover

When you run the `storage aggregate plex show` command after a MetroCluster switchover, the status of `plex0` of the switched over root aggregate is indeterminate and is displayed as `failed`. During this time, the switched over root is not updated. The actual status of this plex can only be determined after the MetroCluster healing phase.

Modifying volumes to set the NVFAIL flag in case of switchover

You can modify a volume so that the NVFAIL flag is set on the volume in the event of a MetroCluster switchover. The NVFAIL flag causes the volume to be fenced off from any modification. This is required for volumes that need to be handled as if committed writes to the volume were lost after the switchover.



In ONTAP versions earlier than 9.0, the NVFAIL flag is used for each switchover. In ONTAP 9.0 and later versions, the unplanned switchover (USO) is used.

Steps

1. Enable MetroCluster configuration to trigger NVFAIL on switchover by setting the `vol -dr-force -nvfail` parameter to "on":

```
vol modify -vserver vserver-name -volume volume-name -dr-force-nvfail on
```

Transitioning from a stretch to a fabric-attached MetroCluster configuration

In a fabric-attached MetroCluster configuration, the nodes are in different locations. This geographical difference increases the disaster protection. To transition from a stretch to a fabric-attached MetroCluster configuration, you must add FC switches and, if necessary, FC-to-SAS bridges to the configuration.

- You must disable automatic switchover on both of the clusters by running the `metrocluster modify -auto-switchover-failure-domain auto-disabled` command.
- You must have shut down the nodes.

This procedure is disruptive.

The MetroCluster configuration must be transitioned on both sites. After upgrading the MetroCluster configuration, you must enable automatic switchover on both the clusters. You also must validate the configuration by running the `metrocluster check run` command.

This procedure gives an overview of the required steps. For detailed steps, you must refer to specific sections in the [Fabric-attached MetroCluster installation and configuration](#). You do not need to do a full installation and configuration.

Steps

1. Prepare for the upgrade by carefully reviewing the "Preparing for the MetroCluster installation" section of the [Fabric-attached MetroCluster installation and configuration](#).
2. Install, cable, and configure the required switches and FC-to-SAS bridges.



You should use the procedures in the section "Cabling a fabric-attached MetroCluster configuration" of the [Fabric-attached MetroCluster installation and configuration](#).

3. Refresh the MetroCluster configuration using the following steps.

Do not use the procedures in the section "Configuring the MetroCluster software in ONTAP" found in the [Fabric-attached MetroCluster installation and configuration](#).

a. Enter advanced privilege mode:

+

```
set -privilege advanced
```

b. Refresh the MetroCluster configuration:

+

```
metrocluster configure -refresh true
```

The following command refreshes the MetroCluster configuration on all the nodes in the DR group that contains controller_A_1:

```
controller_A_1::*> metrocluster configure -refresh true
[Job 009] Job succeeded: Configure is successful.
```

c. Return to admin privilege mode:

+

```
set -privilege admin
```

4. Check the MetroCluster configuration for errors and verify that it is operational.

You should use the procedures in the following sections of the [Fabric-attached MetroCluster installation and configuration](#):

- Checking for MetroCluster configuration errors with Config Advisor
- Verifying local HA operation
- Verifying switchover, healing, and switchback

Where to find additional information

You can learn more about the MetroCluster configuration and operation.

MetroCluster and miscellaneous information

Information	Subject
ONTAP 9 Documentation	<ul style="list-style-type: none">• All MetroCluster guides
	<ul style="list-style-type: none">• A technical overview of the MetroCluster FC configuration and operation.• Best practices for MetroCluster FC configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none">• Fabric-attached MetroCluster architecture• Cabling the configuration• Configuring the FC-to-SAS bridges• Configuring the FC switches• Configuring the MetroCluster in ONTAP

MetroCluster IP installation and configuration: Differences among the ONTAP MetroCluster configurations	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the configuration • Configuring the MetroCluster in ONTAP
MetroCluster management and disaster recovery	<ul style="list-style-type: none"> • Understanding the MetroCluster configuration • Switchover, healing and switchback • Disaster recovery (DR)
Maintain MetroCluster Components	<ul style="list-style-type: none"> • Guidelines for maintenance in a MetroCluster FC configuration • Hardware replacement or upgrade. Firmware upgrade procedures for FC-to-SAS bridges and FC switches • Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Replacing hardware at a disaster recovery site in a fabric-attached or stretch MetroCluster FC configuration • Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration. • Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.
Transition from MetroCluster FC to MetroCluster IP MetroCluster Upgrade and Expansion Guide	<ul style="list-style-type: none"> • Upgrading or refreshing a MetroCluster configuration • Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration • Expanding a MetroCluster configuration by adding additional nodes
MetroCluster Tiebreaker Software installation and configuration	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
Active IQ Unified Manager documentation NetApp Documentation: Product Guides and Resources	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems

Install and configure MetroCluster Tiebreaker

What's new in MetroCluster Tiebreaker support

Enhancements to the MetroCluster Tiebreaker software are provided with each release. Here's what's new in recent releases of MetroCluster Tiebreaker.

Enhancements

ONTAP Tiebreaker version	Enhancements
1.7	<ul style="list-style-type: none">• Bug fixes• Adds support for switchover simulation using the CLI
1.6P1	<ul style="list-style-type: none">• Supporting libraries update• Security enhancements
1.6	<ul style="list-style-type: none">• Improved ease of installation• Supporting libraries update• Security enhancements
1.5	<ul style="list-style-type: none">• Supporting libraries update• Security enhancements
1.4	<ul style="list-style-type: none">• Supporting libraries update

OS support matrix

The following table indicates the supported operating systems for each version of Tiebreaker.

OS for Tiebreaker	1.7	1.6P1	1.6	1.5	1.4
Rocky Linux 9.4	Yes	Yes	No	No	No
Rocky Linux 9.0	No	No	Yes	No	No
Rocky Linux 8.10	Yes	Yes	No	No	No
Red Hat Enterprise Linux (RHEL) 9.6	Yes	Yes	No	No	No

RHEL 9.5	Yes	Yes	No	No	No
RHEL 9.4	Yes	Yes	No	No	No
RHEL 9.3	No	No	No	No	No
RHEL 9.2	Yes	Yes	Yes	No	No
RHEL 9.1	No	No	Yes	No	No
RHEL 9.0	No	No	Yes	No	No
RHEL 8.11 - 9.0	No	No	Yes	No	No
RHEL 8.10	Yes	Yes	Yes	No	No
RHEL 8.9	No	No	Yes	No	No
RHEL 8.8	Yes	Yes	Yes	No	No
RHEL 8.1 - 8.7	No	No	Yes	Yes	Yes
RHEL 7 - 7.9	No	No	No	No	Yes
CentOS 7 - 7.9	No	No	No	No	Yes

Overview of the Tiebreaker software

It is helpful to understand what the NetApp MetroCluster Tiebreaker software is and how it distinguishes between types of failures so that you can monitor your MetroCluster configurations efficiently. You use the Tiebreaker CLI to manage settings and monitor the status and operations of MetroCluster configurations.

Detecting failures with NetApp MetroCluster Tiebreaker software

You need the Tiebreaker software only if you want to monitor two clusters and the connectivity status between them from a third site. The Tiebreaker software resides on a Linux host on the third site and enables each partner in a cluster to distinguish between an ISL failure, when inter-site links are down, from a site failure.

After you install the Tiebreaker software on a Linux host, you can configure the clusters in a MetroCluster configuration to monitor for disaster conditions.

The Tiebreaker software can monitor up to 15 MetroCluster configurations simultaneously. It supports a combination of MetroCluster IP, MetroCluster FC, and stretch MetroCluster configurations.

How the Tiebreaker software detects site failures

The NetApp MetroCluster Tiebreaker software checks the reachability of the nodes in a MetroCluster configuration and the cluster to determine whether a site failure has occurred. The Tiebreaker software also triggers an alert under certain conditions.

Components monitored by the Tiebreaker software

The Tiebreaker software monitors each controller in the MetroCluster configuration by establishing redundant connections through multiple paths to a node management LIF and to the cluster management LIF, both hosted on the IP network.

The Tiebreaker software monitors the following components in the MetroCluster configuration:

- Nodes through local node interfaces
- Cluster through the cluster-designated interfaces
- Surviving cluster to evaluate whether it has connectivity to the disaster site (NV interconnect, storage, and intercluster peering)

When there is a loss of connection between the Tiebreaker software and all of the nodes in the cluster and to the cluster itself, the cluster will be declared as “not reachable” by the Tiebreaker software. It takes around three to five seconds to detect a connection failure. If a cluster is unreachable from the Tiebreaker software, the surviving cluster (the cluster that is still reachable) must indicate that all of the links to the partner cluster are severed before the Tiebreaker software triggers an alert.



All of the links are severed if the surviving cluster can no longer communicate with the cluster at the disaster site through FC (NV interconnect and storage) and intercluster peering.

Failure scenarios during which Tiebreaker software triggers an alert

The Tiebreaker software triggers an alert when the cluster (all of the nodes) at the disaster site is down or unreachable and the cluster at the surviving site indicates the “AllLinksSevered” status.

The Tiebreaker software does not trigger an alert (or the alert is vetoed) in the following scenarios:

- In an eight-node MetroCluster configuration, if one HA pair at the disaster site is down
- In a cluster with all of the nodes at the disaster site down, one HA pair at the surviving site down, and the cluster at the surviving site indicates the “AllLinksSevered” status

The Tiebreaker software triggers an alert, but ONTAP vetoes that alert. In this situation, a manual switchover is also vetoed

- Any scenario in which the Tiebreaker software can either reach at least one node or the cluster interface at the disaster site, or the surviving site still can reach either node at the disaster site through either FC (NV interconnect and storage) or intercluster peering

Related information

[Risks and limitations of using MetroCluster Tiebreaker in active mode](#)

How the Tiebreaker software detects intersite connectivity failures

The MetroCluster Tiebreaker software alerts you if all connectivity between the sites is lost.

Types of network paths

Depending on the configuration, there are three types of network paths between the two clusters in a MetroCluster configuration:

- **FC network (present in fabric-attached MetroCluster configurations)**

This type of network is composed of two redundant FC switch fabrics. Each switch fabric has two FC switches, with one switch of each switch fabric co-located with a cluster. Each cluster has two FC switches, one from each switch fabric. All of the nodes have FC (NV interconnect and FCP initiator) connectivity to each of the co-located FC switches. Data is replicated from cluster to cluster over the ISL.

- **Intercluster peering network**

This type of network is composed of a redundant IP network path between the two clusters. The cluster peering network provides the connectivity that is required to mirror the storage virtual machine (SVM) configuration. The configuration of all of the SVMs on one cluster is mirrored by the partner cluster.

- **IP network (present in MetroCluster IP configurations)**

This type of network is composed of two redundant IP switch networks. Each network has two IP switches, with one switch of each switch fabric co-located with a cluster. Each cluster has two IP switches, one from each switch fabric. All of the nodes have connectivity to each of the co-located FC switches. Data is replicated from cluster to cluster over the ISL.

Monitoring intersite connectivity

The Tiebreaker software regularly retrieves the status of intersite connectivity from the nodes. If NV interconnect connectivity is lost and the intercluster peering does not respond to pings, then the clusters assume that the sites are isolated and the Tiebreaker software triggers an alert as “AllLinksSevered”. If a cluster identifies the “AllLinksSevered” status and the other cluster is not reachable through the network, then the Tiebreaker software triggers an alert as “disaster”.

How different disaster types affect Tiebreaker software detection time

For better disaster recovery planning, the MetroCluster Tiebreaker software takes some time in detecting a disaster. This time spent is the “disaster detection time”. The MetroCluster Tiebreaker software detects the site disaster within 30 seconds from the time of occurrence of the disaster and triggers the disaster recovery operation to notify you about the disaster.

The detection time also depends on the type of disaster and might exceed 30 seconds in some scenarios, mostly known as “rolling disasters”. The main types of rolling disaster are as follows:

- Power loss
- Panic
- Halt or reboot
- Loss of FC switches at the disaster site

Power loss

The Tiebreaker software immediately triggers an alert when the node stops operating. When there is a power loss, all connections and updates, such as intercluster peering, NV interconnect, and MailBox disk, stop. The time taken between the cluster becoming unreachable, the detection of the disaster, and the trigger, including

the default silent time of 5 seconds, should not exceed 30 seconds.

Panic

In MetroCluster FC configurations, the Tiebreaker software triggers an alert when the NV interconnect connection between the sites is down and the surviving site indicates the “AllLinksSevered” status. This only happens after the coredump process is complete. In this scenario, the time taken between the cluster becoming unreachable and the detection of a disaster might be longer or approximately equal to the time taken for the coredump process. In many cases, the detection time is more than 30 seconds.

If a node stops operating but does not generate a file for the coredump process, then the detection time should not be longer than 30 seconds.

In MetroCluster IP configurations, the NV stops communicating and the surviving site is not aware of the coredump process.

Halt or reboot

The Tiebreaker software triggers an alert only when the node is down and the surviving site indicates the “AllLinksSevered” status. The time taken between the cluster becoming unreachable and the detection of a disaster might be longer than 30 seconds. In this scenario, the time taken to detect a disaster depends on how long it takes for the nodes at the disaster site to be shut down.

Loss of FC switches at the disaster site (fabric-attached MetroCluster configuration)

The Tiebreaker software triggers an alert when a node stops operating. If FC switches are lost, then the node tries to recover the path to a disk for about 30 seconds. During this time, the node is up and responding on the peering network. When both of the FC switches are down and the path to a disk cannot be recovered, the node produces a MultiDiskFailure error and halts. The time taken between the FC switch failure and the number of times the nodes produced MultiDiskFailure errors is about 30 seconds longer. This additional 30 seconds must be added to the disaster detection time.

About the Tiebreaker CLI and man pages

The Tiebreaker CLI provides commands that enable you to remotely configure the Tiebreaker software and monitor the MetroCluster configurations.

The CLI command prompt is represented as NetApp MetroCluster Tiebreaker::>.

The man pages are available in the CLI by entering the applicable command name at the prompt.

Install the Tiebreaker software

Tiebreaker installation workflow

The Tiebreaker software provides monitoring capabilities for a clustered storage environment. It also sends SNMP notifications in the event of node connectivity issues and site disasters.

About this workflow

You can use this workflow to install or upgrade the Tiebreaker software.

1

Prepare to install the Tiebreaker software

Before you install and configure the Tiebreaker software, verify that your system meets certain requirements.

2

Secure the installation

For configurations running MetroCluster Tiebreaker 1.5 and later, you can secure and harden the host OS and the database.

3

Install the Tiebreaker software package

Perform a new installation or upgrade of the Tiebreaker software. The installation procedure you follow depends on the version of Tiebreaker you want to install.

Prepare to install the Tiebreaker software

Before you install and configure the Tiebreaker software you should verify that your system meets certain requirements.

Software requirements

You must meet the following software requirements depending on the version of Tiebreaker you are installing.

ONTAP Tiebreaker version	Supported ONTAP versions	Supported Linux versions	Java/MariaDB requirements
1.7	ONTAP 9.12.1 and later	Refer to the OS Support Matrix for details.	None. The dependencies are bundled with the installation.
1.6P1	ONTAP 9.12.1 and later	Refer to the OS Support Matrix for details.	None. The dependencies are bundled with the installation.
1.6	ONTAP 9.12.1 and later	Refer to the OS Support Matrix for details.	None. The dependencies are bundled with the installation.
1.5	ONTAP 9.8 to ONTAP 9.14.1	<ul style="list-style-type: none"> Red Hat Enterprise Linux 8.1 to 8.7 	With Red Hat Enterprise Linux 8.1 to 8.7: <ul style="list-style-type: none"> MariaDB 10.x (use the default version that is installed using "yum install mariadb-server.x86_64") OpenJDK 17, 18, or 19

1.4	ONTAP 9.1 to ONTAP 9.9.1	<ul style="list-style-type: none"> • Red Hat Enterprise Linux 8.1 to 8.7 • Red Hat Enterprise Linux 7 to 7.9 • CentOS 7 to 7.9 64-bit 	<p>With CentOS:</p> <ul style="list-style-type: none"> • MariaDB 5.5.52.x/MySQL Server 5.6x • 4 GB RAM • Open JRE 8 <p>With Red Hat Enterprise Linux 8.1 to 8.7:</p> <ul style="list-style-type: none"> • MariaDB 10.x (use the default version that is installed using "yum install mariadb-server.x86_64") • JRE 8
-----	-----------------------------	--	---

Additional requirements

You must be aware of the following additional requirements:

- The Tiebreaker software is installed on a third site, which allows the software to distinguish between an inter-switch link (ISL) failure (when inter-site links are down) and a site failure. Your host system must meet certain requirements before you can install or upgrade the Tiebreaker software to monitor the MetroCluster configuration.
- You must have "root" privileges to install MetroCluster Tiebreaker software and the dependant packages.
- You can only use one MetroCluster Tiebreaker monitor per MetroCluster configuration to avoid any conflict with multiple Tiebreaker monitors.
- When selecting the Network Time Protocol (NTP) source for the Tiebreaker software, you must use a local NTP source. The Tiebreaker software should not use the same source as the MetroCluster sites that the Tiebreaker software monitors.
- Disk capacity: 8 GB
- Firewall:
 - Direct access for setting up AutoSupport messages
 - SSH (port 22/TCP), HTTPS (port 443/TCP), and ping (ICMP)

Secure the Tiebreaker host and database installation

For configurations running MetroCluster Tiebreaker 1.5 and later, you can secure and harden the host OS and the database.

Secure the host

The following guidelines show you how to secure the host where the Tiebreaker software is installed.

User management recommendations

- Limit access of the "root" user.
 - You can use users that are able to elevate to root access to install and administer the Tiebreaker software.

- You can use users that are not able to elevate to root access to administer Tiebreaker software.
- During installation, you must create a group named "mcctbgrp". The host root user and the user created during the installation must both be members. Only members of this group can fully administer the Tiebreaker software.



Users who are not members of this group cannot access the Tiebreaker software or CLI. You can create additional users on the host and make them members of the group. These additional members cannot fully administer the Tiebreaker software. They have ReadOnly access and cannot add, change, or delete monitors.

- Do not run Tiebreaker as a root user. Use a dedicated, unprivileged service account to run Tiebreaker.
- Change the default community string in the "/etc/snmp/snmpd.conf" file.
- Allow minimal write privileges. The unprivileged Tiebreaker service account should not have access to overwrite its executable binary or any configuration files. Only directories and files for local Tiebreaker storage (eg., for integrated backend storage) or audit logs should be writable by the Tiebreaker user.
- Do not permit anonymous users.
 - Set AllowTcpForwarding to "no" or use the Match directive to restrict anonymous users.

Related information

- [Red Hat Enterprise Linux 8 product documentation](#)
- [Red Hat Enterprise Linux 9 product documentation](#)
- [Rocky Linux product documentation](#)

Baseline host security recommendations

- Use disk encryption
 - You can enable disk encryption. This can be FullDiskEncryption (hardware), or encryption provided by the HostOS (software), or by the SVM host.
- Disable unused services that allow incoming connections. You can disable any service that isn't in use. The Tiebreaker software does not require a service for incoming connections because all connections from the Tiebreaker installation are outgoing. The services that might be enabled by default and can be disabled are:
 - HTTP/HTTPS server
 - FTP server
 - Telnet, RSH, rlogin
 - NFS, CIFS, and other protocol access
 - RDP (RemoteDesktopProtocol), X11 Server, VNC or other remote "desktop" service providers.



You must leave either serial console access (if supported) or at least one protocol enabled to administer the host remotely. If you disable all protocols, then you require physical access to the host for administration.

- Secure the host using FIPS
 - You can install the host OS in FIPS-compliant mode and then install Tiebreaker.



OpenJDK 19 checks on startup whether the host is installed in FIPS mode. No manual changes should be required.

- If you secure the host, you must ensure that the host is able to boot without user intervention. If user intervention is required, Tiebreaker functionality might not be available if the host unexpectedly reboots. If this occurs, Tiebreaker functionality is only available after the manual intervention and when the host is fully booted.
- Disable Shell Command History.
- Upgrade frequently. Tiebreaker is actively developed, and updating frequently is important to incorporate security fixes and any changes in default settings such as key lengths or cipher suites.
- Subscribe to the HashiCorp Announcement mailing list to receive announcements of new releases and visit the Tiebreaker CHANGELOG for details on recent updates for new releases.
- Use the correct file permissions. Always ensure appropriate permissions are applied to files before starting the Tiebreaker software, especially those containing sensitive information.
- Multifactor authentication (MFA) enhances your organization's security by requiring administrators to identify themselves by using more than a username and password. While important, usernames and passwords are vulnerable to brute force attacks and can be stolen by third parties.
 - Red Hat Enterprise Linux 8 provides MFA that requires users to provide more than one piece of information to authenticate successfully to an account or Linux host. The additional information might be a one-time password sent to your cell phone via SMS or credentials from an app like Google Authenticator, Twilio Authy, or FreeOTP.

Related information

- [Red Hat Enterprise Linux 8 product documentation](#)
- [Red Hat Enterprise Linux 9 product documentation](#)
- [Rocky Linux product documentation](#)

Secure the database installation

The following guidelines show how to secure and harden the MariaDB 10.x database installation.

- Limit the access of the "root" user.
 - Tiebreaker uses a dedicated account. The account and tables for storing (configuration) data is created during the installation of Tiebreaker. The only time elevated access to the database is required is during installation.
- During installation the following access and privileges are required:
 - The ability to create a database and tables
 - The ability to create global options
 - The ability to create a database user and set the password
 - The ability to associate the database user with the database and tables and assign access rights



The user account that you specify during the Tiebreaker installation must have all these privileges. Using multiple user accounts for the different tasks is not supported.

- Use encryption of the database
 - Data-at-rest encryption is supported. [Learn more about data-at-rest encryption](#)

- Data in flight is not encrypted. Data in flight uses a local "socks" file connection.
- FIPS compliancy for MariaDB — you do not need to enable FIPS compliancy on the database. Installation of the host in FIPS-compliant mode is sufficient.

[Learn about MySQL Enterprise Transparent Data Encryption \(TDE\)](#)



The encryption settings must be enabled before installation of the Tiebreaker software.

Related information

- Database user management

[Access Control and Account Management](#)

- Secure the database

[Making MySQL Secure Against Attackers](#)

[Securing MariaDB](#)

- Secure the Vault installation

[Production hardening](#)

Install the Tiebreaker software package

Choose your installation procedure

The Tiebreaker installation procedure you follow depends on the version of Tiebreaker you are installing.

Tiebreaker version	Go to...
Tiebreaker 1.7	Install Tiebreaker 1.7
Tiebreaker 1.6 or 1.6P1	Install Tiebreaker 1.6 or 1.6P1
Tiebreaker 1.5	Install Tiebreaker 1.5
Tiebreaker 1.4	Install Tiebreaker 1.4

Install MetroCluster Tiebreaker 1.7

Install or upgrade to Tiebreaker 1.7 on your Linux host to monitor MetroCluster configurations.

About this task

- Your storage system must be running ONTAP 9.12.1 or later.
- You can only upgrade to Tiebreaker 1.7 from Tiebreaker 1.6P1. Refer to [Install Tiebreaker 1.6 or 1.6P1](#).

- You can install MetroCluster Tiebreaker as a non-root user with sufficient administrative privileges to perform the Tiebreaker installation, create tables and users, and set the user password.

Steps

1. Download the MetroCluster Tiebreaker 1.7 software.

[MetroCluster Tiebreaker \(Downloads\) - NetApp Support Site](#)

2. Log in to the host as the root user.
3. If you are upgrading, check which version of Tiebreaker you are running:

The following example shows Tiebreaker 1.6P1

```
[root@mcctb ~] # netapp-metrocluster-tiebreaker-software-cli
NetApp MetroCluster Tiebreaker :> version show
NetApp MetroCluster Tiebreaker 1.6P1: Sun Mar 13 09:59:02 IST 2022
NetApp MetroCluster Tiebreaker :> exit
```

4. Install or upgrade the Tiebreaker software.

Install Tiebreaker 1.7

Use the following steps for a new installation of Tiebreaker 1.7.

Steps

1. Run the following command at the `[root@mcctb ~] #` prompt to begin the installation:

```
sh MetroClusterTiebreakerInstall-1.7
```

The system displays the following output for a successful installation:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
Enter unix user account to use for the installation:
mcctbadminuser
Unix user account "mcctbadminuser" doesn't exist. Do you wish
to create "mcctbadminuser" user account? [Y/N]: y
useradd: warning: the home directory already exists.
Not copying any file from skel directory into it.
Creating mailbox file: File exists
Unix account "mcctbadminuser" created.
Changing password for user mcctbadminuser.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
MetroCluster Tiebreaker requires unix user account
"mcctbadminuser" to be added to the group "mcctbgrp" for admin
access.
Do you wish to add ? [Y/N]: y
Unix user account "mcctbadminuser" added to "mcctbgrp".
Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
Starting vault server...
==> Vault server configuration:

        Api Address: <api_address>
                Cgo: disabled
        Cluster Address: <cluster_address>
        Environment Variables: BASH_FUNC_which%%,
        DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
        HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
        LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
        SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,
        VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, _,
        vault_Addr, which_declare
        Go Version: go1.20.5
        Listener 1: tcp (addr: "0.0.0.0:8200", cluster
        address: "0.0.0.0:8201", max_request_duration: "1m30s",
        max_request_size: "33554432", tls: "enabled")
        Log Level:
```

```

                Mlock: supported: true, enabled: true
                Recovery Mode: false
                Storage: file
                Version: Vault v1.14.0, built 2023-06-
19T11:40:23Z
                Version Sha:
13a649f860186dffe3f3a4459814d87191efc321

==> Vault server started! Log data will stream in below:

2023-11-23T15:14:28.532+0530 [INFO] proxy environment:
http_proxy="" https_proxy="" no_proxy=""
2023-11-23T15:14:28.577+0530 [INFO] core: Initializing
version history cache for core
2023-11-23T15:14:38.552+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:14:38.552+0530 [INFO] core: seal configuration
missing, not initialized
2023-11-23T15:14:38.554+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:14:38.555+0530 [INFO] core: security barrier
initialized: stored=1 shares=5 threshold=3
2023-11-23T15:14:38.556+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:14:38.577+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:14:38.577+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:14:38.577+0530 [INFO] core: no mounts; adding
default mount table
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:14:38.581+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-11-23T15:14:38.581+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:14:38.581+0530 [INFO] core: restoring leases
2023-11-23T15:14:38.582+0530 [INFO] expiration: lease restore
```

```
complete
2023-11-23T15:14:38.582+0530 [INFO] identity: entities
restored
2023-11-23T15:14:38.582+0530 [INFO] identity: groups restored
2023-11-23T15:14:38.583+0530 [INFO] core: Recorded vault
version: vault version=1.14.0 upgrade time="2023-11-23
09:44:38.582881162 +0000 UTC" build date=2023-06-19T11:40:23Z
2023-11-23T15:14:38.583+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:14:38.998+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:14:38.999+0530 [INFO] core: root token
generated
2023-11-23T15:14:38.999+0530 [INFO] core: pre-seal teardown
starting
2023-11-23T15:14:38.999+0530 [INFO] rollback: stopping
rollback manager
2023-11-23T15:14:38.999+0530 [INFO] core: pre-seal teardown
complete
2023-11-23T15:14:39.311+0530 [INFO] core.cluster-
listener.tcp: starting listener: listener_address=0.0.0.0:8201
2023-11-23T15:14:39.311+0530 [INFO] core.cluster-listener:
serving cluster requests: cluster_listen_address=[:]:8201
2023-11-23T15:14:39.312+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:14:39.312+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:14:39.312+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:14:39.314+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-11-23T15:14:39.314+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:14:39.314+0530 [INFO] core: restoring leases
2023-11-23T15:14:39.314+0530 [INFO] identity: entities
restored
```

```
2023-11-23T15:14:39.314+0530 [INFO] expiration: lease restore
complete
2023-11-23T15:14:39.314+0530 [INFO] identity: groups restored
2023-11-23T15:14:39.315+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:14:39.316+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:14:39.316+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-11-23T15:14:39.795+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-11-23T15:14:39.885+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Installing the NetApp-MetroCluster-Tiebreaker-Software-1.7-
1.x86_64.rpm
Preparing... #
##### # [100%]

Updating / installing...

1:NetApp-MetroCluster-Tiebreaker-So#
##### # [100%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
```

opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok

```
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/
```

```
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
Created symlink /etc/systemd/system/multi-
user.target.wants/netapp-metrocluster-tiebreaker-
software.service → /etc/systemd/system/netapp-metrocluster-
tiebreaker-software.service.
```

```
Attempting to start NetApp MetroCluster Tiebreaker software
services
```

```
Started NetApp MetroCluster Tiebreaker software services
```

```
Successfully installed NetApp MetroCluster Tiebreaker software
version 1.7.
```

Upgrade 1.6P1 to 1.7

Use the following steps to upgrade the Tiebreaker 1.6P1 software version to Tiebreaker 1.7.

Steps

1. Run the following command at the [root@mcctb ~] # prompt to upgrade the software:

```
sh MetroClusterTiebreakerInstall-1.7
```

The system displays the following output for a successful upgrade:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
NetApp-MetroCluster-Tiebreaker-Software-1.6P1-1.x86_64
Upgrading... to NetApp-MetroCluster-Tiebreaker-Software-1.7-
1.x86_64.rpm
Preparing...
##### [100%]
Updating / installing...
  1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.19.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
```

```
opt/netapp/mcctb/lib/common/jakarta.mail-2.0.1.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/
```

```
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
```

```
to version 1.7.  
Cleaning up / removing...  
2:NetApp-MetroCluster-Tiebreaker-  
So##### [100%]
```



If your upgrade fails with an error due to missing certificates, refer to [Import certificates](#). After your certificates are imported, you can try to upgrade again.

Install Tiebreaker 1.6 or 1.6P1

Perform a new installation or upgrade to Tiebreaker 1.6 or Tiebreaker 1.6P1 on your host Linux operating system to monitor MetroCluster configurations.

About this task

- Your storage system must be running ONTAP 9.12.1 or later.
- You can install MetroCluster Tiebreaker as a non-root user with sufficient administrative privileges to perform the Tiebreaker installation, create tables and users, and set the user password.

Install or upgrade to Tiebreaker 1.6P1

You can install Tiebreaker 1.6P1 or upgrade to Tiebreaker 1.6P1 from Tiebreaker 1.6, 1.5, or 1.4.

Steps

1. Download the MetroCluster Tiebreaker 1.6P1 software.

[MetroCluster Tiebreaker \(Downloads\) - NetApp Support Site](#)

2. Log in to the host as the root user.
3. If you are performing an upgrade, verify the version of Tiebreaker that you are running:

The following example shows Tiebreaker 1.5.

```
[root@mcctb ~] # netapp-metrocluster-tiebreaker-software-cli  
NetApp MetroCluster Tiebreaker :> version show  
NetApp MetroCluster Tiebreaker 1.5: Sun Mar 13 09:59:02 IST 2022  
NetApp MetroCluster Tiebreaker :> exit
```

4. Install or upgrade the Tiebreaker software.

Install Tiebreaker 1.6P1

Use the following steps for a new installation of Tiebreaker 1.6P1.

Steps

1. Run the following command at the `[root@mcctb ~] #` prompt to begin the installation:

```
sh MetroClusterTiebreakerInstall-1.6P1
```

The system displays the following output for a successful installation:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
Enter unix user account to use for the installation:
mcctbadminuser
Unix user account "mcctbadminuser" doesn't exist. Do you wish
to create "mcctbadminuser" user account? [Y/N]: y
useradd: warning: the home directory already exists.
Not copying any file from skel directory into it.
Creating mailbox file: File exists
Unix account "mcctbadminuser" created.
Changing password for user mcctbadminuser.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
MetroCluster Tiebreaker requires unix user account
"mcctbadminuser" to be added to the group "mcctbgrp" for admin
access.
Do you wish to add ? [Y/N]: y
Unix user account "mcctbadminuser" added to "mcctbgrp".
Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
Starting vault server...
==> Vault server configuration:

      Api Address: <api_address>
          Cgo: disabled
      Cluster Address: <cluster_address>
  Environment Variables: BASH_FUNC_which%%,
  DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
  HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
  LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
  SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,
  VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, _,
  vault_Addr, which_declare
      Go Version: go1.20.5
      Listener 1: tcp (addr: "0.0.0.0:8200", cluster
  address: "0.0.0.0:8201", max_request_duration: "1m30s",
  max_request_size: "33554432", tls: "enabled")
      Log Level:
```

```
        Mlock: supported: true, enabled: true
        Recovery Mode: false
        Storage: file
        Version: Vault v1.14.0, built 2023-06-
19T11:40:23Z
        Version Sha:
13a649f860186dffe3f3a4459814d87191efc321

==> Vault server started! Log data will stream in below:

2023-11-23T15:14:28.532+0530 [INFO] proxy environment:
http_proxy="" https_proxy="" no_proxy=""
2023-11-23T15:14:28.577+0530 [INFO] core: Initializing
version history cache for core
2023-11-23T15:14:38.552+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:14:38.552+0530 [INFO] core: seal configuration
missing, not initialized
2023-11-23T15:14:38.554+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:14:38.555+0530 [INFO] core: security barrier
initialized: stored=1 shares=5 threshold=3
2023-11-23T15:14:38.556+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:14:38.577+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:14:38.577+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:14:38.577+0530 [INFO] core: no mounts; adding
default mount table
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:14:38.581+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-11-23T15:14:38.581+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:14:38.581+0530 [INFO] core: restoring leases
2023-11-23T15:14:38.582+0530 [INFO] expiration: lease restore
```

```
complete
2023-11-23T15:14:38.582+0530 [INFO] identity: entities
restored
2023-11-23T15:14:38.582+0530 [INFO] identity: groups restored
2023-11-23T15:14:38.583+0530 [INFO] core: Recorded vault
version: vault version=1.14.0 upgrade time="2023-11-23
09:44:38.582881162 +0000 UTC" build date=2023-06-19T11:40:23Z
2023-11-23T15:14:38.583+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:14:38.998+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:14:38.999+0530 [INFO] core: root token
generated
2023-11-23T15:14:38.999+0530 [INFO] core: pre-seal teardown
starting
2023-11-23T15:14:38.999+0530 [INFO] rollback: stopping
rollback manager
2023-11-23T15:14:38.999+0530 [INFO] core: pre-seal teardown
complete
2023-11-23T15:14:39.311+0530 [INFO] core.cluster-
listener.tcp: starting listener: listener_address=0.0.0.0:8201
2023-11-23T15:14:39.311+0530 [INFO] core.cluster-listener:
serving cluster requests: cluster_listen_address=[:]:8201
2023-11-23T15:14:39.312+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:14:39.312+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:14:39.312+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:14:39.314+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-11-23T15:14:39.314+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:14:39.314+0530 [INFO] core: restoring leases
2023-11-23T15:14:39.314+0530 [INFO] identity: entities
restored
```

```
2023-11-23T15:14:39.314+0530 [INFO] expiration: lease restore
complete
2023-11-23T15:14:39.314+0530 [INFO] identity: groups restored
2023-11-23T15:14:39.315+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:14:39.316+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:14:39.316+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-11-23T15:14:39.795+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-11-23T15:14:39.885+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Installing the NetApp-MetroCluster-Tiebreaker-Software-1.6P1-
1.x86_64.rpm
Preparing... #
##### # [100%]

Updating / installing...

1:NetApp-MetroCluster-Tiebreaker-So#
##### # [100%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
```

opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok

```
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/
```

```
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
Created symlink /etc/systemd/system/multi-
user.target.wants/netapp-metrocluster-tiebreaker-
software.service → /etc/systemd/system/netapp-metrocluster-
tiebreaker-software.service.
```

```
Attempting to start NetApp MetroCluster Tiebreaker software
services
```

```
Started NetApp MetroCluster Tiebreaker software services
```

```
Successfully installed NetApp MetroCluster Tiebreaker software
version 1.6P1.
```

Upgrade 1.6 to 1.6P1

Use the following steps to upgrade the Tiebreaker 1.6 software version to Tiebreaker 1.6P1.



After you upgrade to Tiebreaker 1.6P1 from 1.6, you remove the existing monitors and re-add the MetroCluster configuration for monitoring.

Steps

1. Run the following command at the [root@mcctb ~] # prompt to upgrade the software:

```
sh MetroClusterTiebreakerInstall-1.6P1
```

The system displays the following output for a successful upgrade:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
NetApp-MetroCluster-Tiebreaker-Software-1.6P1-1.x86_64
Error making API request.

URL: GET
https://127.0.0.1:8200/v1/sys/internal/ui/mounts/mcctb/data/db
Code: 403. Errors:

* permission denied
Upgrading to NetApp-MetroCluster-Tiebreaker-Software-1.6P1-
1.x86_64.rpm
Preparing...
##### [100%]
Updating / installing...
  1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.19.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
```

opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.mail-2.0.1.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok

```
/
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
to version 1.6P1.
Cleaning up / removing...
      2:NetApp-MetroCluster-Tiebreaker-
So##### [100%]
```

2. Remove and re-add the MetroCluster configuration by following the steps in [Configure the Tiebreaker software](#).

Upgrade 1.5 to 1.6P1

Use the following steps to upgrade the Tiebreaker 1.5 software version to Tiebreaker 1.6P1.

Steps

1. Run the following command at the [root@mcctb ~] # prompt to upgrade the software:

```
sh MetroClusterTiebreakerInstall-1.6P1
```

The system displays the following output for a successful upgrade:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok

Enter database user name : root

Please enter database password for root
Enter password:

Password updated successfully in the database.

Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
==> Vault shutdown triggered
2023-07-21T00:30:22.335+0530 [INFO] core: marked as sealed
2023-07-21T00:30:22.335+0530 [INFO] core: pre-seal teardown
starting
2023-07-21T00:30:22.335+0530 [INFO] rollback: stopping
rollback manager
2023-07-21T00:30:22.335+0530 [INFO] core: pre-seal teardown
complete
2023-07-21T00:30:22.335+0530 [INFO] core: stopping cluster
listeners
2023-07-21T00:30:22.335+0530 [INFO] core.cluster-listener:
forwarding rpc listeners stopped
2023-07-21T00:30:22.375+0530 [INFO] core.cluster-listener:
rpc listeners successfully shut down
2023-07-21T00:30:22.375+0530 [INFO] core: cluster listeners
successfully shut down
2023-07-21T00:30:22.376+0530 [INFO] core: vault is sealed
Starting vault server...
==> Vault server configuration:

        Api Address: <api_address>
                Cgo: disabled
        Cluster Address: <cluster_address>
        Environment Variables: BASH_FUNC_which%%,
        DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
        HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
        LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
```

```
SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,  
VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, _,  
vault_Addr, which_declare
```

```
Go Version: go1.20.5
```

```
Listener 1: tcp (addr: "0.0.0.0:8200", cluster  
address: "0.0.0.0:8201", max_request_duration: "1m30s",  
max_request_size: "33554432", tls: "enabled")
```

```
Log Level:
```

```
Mlock: supported: true, enabled: true
```

```
Recovery Mode: false
```

```
Storage: file
```

```
Version: Vault v1.14.0, built 2023-06-
```

```
19T11:40:23Z
```

```
Version Sha:
```

```
13a649f860186dffe3f3a4459814d87191efc321
```

```
==> Vault server started! Log data will stream in below:
```

```
2023-07-21T00:30:33.065+0530 [INFO] proxy environment:  
http_proxy="" https_proxy="" no_proxy=""  
2023-07-21T00:30:33.098+0530 [INFO] core: Initializing  
version history cache for core  
2023-07-21T00:30:43.092+0530 [INFO] core: security barrier  
not initialized  
2023-07-21T00:30:43.092+0530 [INFO] core: seal configuration  
missing, not initialized  
2023-07-21T00:30:43.094+0530 [INFO] core: security barrier  
not initialized  
2023-07-21T00:30:43.096+0530 [INFO] core: security barrier  
initialized: stored=1 shares=5 threshold=3  
2023-07-21T00:30:43.098+0530 [INFO] core: post-unseal setup  
starting  
2023-07-21T00:30:43.124+0530 [INFO] core: loaded wrapping  
token key  
2023-07-21T00:30:43.124+0530 [INFO] core: successfully setup  
plugin catalog: plugin-directory=""  
2023-07-21T00:30:43.124+0530 [INFO] core: no mounts; adding  
default mount table  
2023-07-21T00:30:43.125+0530 [INFO] core: successfully  
mounted: type=cubbyhole version="v1.14.0+builtin.vault"  
path=cubbyhole/ namespace="ID: root. Path: "  
2023-07-21T00:30:43.126+0530 [INFO] core: successfully  
mounted: type=system version="v1.14.0+builtin.vault" path=sys/  
namespace="ID: root. Path: "  
2023-07-21T00:30:43.126+0530 [INFO] core: successfully  
mounted: type=identity version="v1.14.0+builtin.vault"
```

```
path=identity/ namespace="ID: root. Path: "
2023-07-21T00:30:43.129+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-07-21T00:30:43.130+0530 [INFO] rollback: starting
rollback manager
2023-07-21T00:30:43.130+0530 [INFO] core: restoring leases
2023-07-21T00:30:43.130+0530 [INFO] identity: entities
restored
2023-07-21T00:30:43.130+0530 [INFO] identity: groups restored
2023-07-21T00:30:43.131+0530 [INFO] core: usage gauge
collection is disabled
2023-07-21T00:30:43.131+0530 [INFO] expiration: lease restore
complete
2023-07-21T00:30:43.131+0530 [INFO] core: Recorded vault
version: vault version=1.14.0 upgrade time="2023-07-20
19:00:43.131158543 +0000 UTC" build date=2023-06-19T11:40:23Z
2023-07-21T00:30:43.371+0530 [INFO] core: post-unseal setup
complete
2023-07-21T00:30:43.371+0530 [INFO] core: root token
generated
2023-07-21T00:30:43.371+0530 [INFO] core: pre-seal teardown
starting
2023-07-21T00:30:43.371+0530 [INFO] rollback: stopping
rollback manager
2023-07-21T00:30:43.372+0530 [INFO] core: pre-seal teardown
complete
2023-07-21T00:30:43.694+0530 [INFO] core.cluster-
listener.tcp: starting listener: listener_address=0.0.0.0:8201
2023-07-21T00:30:43.695+0530 [INFO] core.cluster-listener:
serving cluster requests: cluster_listen_address=[:]:8201
2023-07-21T00:30:43.695+0530 [INFO] core: post-unseal setup
starting
2023-07-21T00:30:43.696+0530 [INFO] core: loaded wrapping
token key
2023-07-21T00:30:43.696+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-07-21T00:30:43.697+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-07-21T00:30:43.698+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-07-21T00:30:43.698+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
```

```
2023-07-21T00:30:43.701+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-07-21T00:30:43.701+0530 [INFO] rollback: starting
rollback manager
2023-07-21T00:30:43.702+0530 [INFO] core: restoring leases
2023-07-21T00:30:43.702+0530 [INFO] identity: entities
restored
2023-07-21T00:30:43.702+0530 [INFO] expiration: lease restore
complete
2023-07-21T00:30:43.702+0530 [INFO] identity: groups restored
2023-07-21T00:30:43.702+0530 [INFO] core: usage gauge
collection is disabled
2023-07-21T00:30:43.703+0530 [INFO] core: post-unseal setup
complete
2023-07-21T00:30:43.703+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-07-21T00:30:44.226+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-07-21T00:30:44.315+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Upgrading to NetApp-MetroCluster-Tiebreaker-Software-1.6P1-
1.x86_64.rpm
Preparing...
##### [100%]
Updating / installing...
 1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
```

opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbc2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok

```
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/
```

```
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
```

```
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
to version 1.6P1.
Cleaning up / removing...
  2:NetApp-MetroCluster-Tiebreaker-
So##### [100%]
```

Upgrade 1.4 to 1.6P1

Use the following steps to upgrade the Tiebreaker 1.4 software version to Tiebreaker 1.6P1.

Steps

1. Run the following command at the [root@mcctb ~] # prompt to upgrade the software:

```
sh MetroClusterTiebreakerInstall-1.6P1
```

The system displays the following output for a successful upgrade:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
Enter unix user account to use for the installation:
mcctbuseradmin1
Unix user account "mcctbuseradmin1" doesn't exist. Do you wish
to create "mcctbuseradmin1" user account? [Y/N]: y
Unix account "mcctbuseradmin1" created.
Changing password for user mcctbuseradmin1.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.

Enter database user name : root

Please enter database password for root
Enter password:

Password updated successfully in the database.

MetroCluster Tiebreaker requires unix user account
"mcctbuseradmin1" to be added to the group "mcctbgrp" for
admin access.
Do you wish to add ? [Y/N]: y
Unix user account "mcctbuseradmin1" added to "mcctbgrp".
Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
Starting vault server...
==> Vault server configuration:

      Api Address: <api_address>
            Cgo: disabled
      Cluster Address: <cluster_address>
  Environment Variables: BASH_FUNC_which%,
DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,
VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, __,
vault_Addr, which_declare
```

```
Go Version: go1.20.5
Listener 1: tcp (addr: "0.0.0.0:8200", cluster
address: "0.0.0.0:8201", max_request_duration: "1m30s",
max_request_size: "33554432", tls: "enabled")
Log Level:
Mlock: supported: true, enabled: true
Recovery Mode: false
Storage: file
Version: Vault v1.14.0, built 2023-06-
19T11:40:23Z
Version Sha:
13a649f860186dffe3f3a4459814d87191efc321
```

==> Vault server started! Log data will stream in below:

```
2023-11-23T15:58:10.400+0530 [INFO] proxy environment:
http_proxy="" https_proxy="" no_proxy=""
2023-11-23T15:58:10.432+0530 [INFO] core: Initializing
version history cache for core
2023-11-23T15:58:20.422+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:58:20.422+0530 [INFO] core: seal configuration
missing, not initialized
2023-11-23T15:58:20.424+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:58:20.425+0530 [INFO] core: security barrier
initialized: stored=1 shares=5 threshold=3
2023-11-23T15:58:20.427+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:58:20.448+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:58:20.448+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:58:20.448+0530 [INFO] core: no mounts; adding
default mount table
2023-11-23T15:58:20.449+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:58:20.449+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:58:20.449+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:58:20.451+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
```

```
path=token/ namespace="ID: root. Path: "  
2023-11-23T15:58:20.452+0530 [INFO] rollback: starting  
rollback manager  
2023-11-23T15:58:20.452+0530 [INFO] core: restoring leases  
2023-11-23T15:58:20.453+0530 [INFO] identity: entities  
restored  
2023-11-23T15:58:20.453+0530 [INFO] identity: groups restored  
2023-11-23T15:58:20.453+0530 [INFO] expiration: lease restore  
complete  
2023-11-23T15:58:20.453+0530 [INFO] core: usage gauge  
collection is disabled  
2023-11-23T15:58:20.453+0530 [INFO] core: Recorded vault  
version: vault version=1.14.0 upgrade time="2023-11-23  
10:28:20.453481904 +0000 UTC" build date=2023-06-19T11:40:23Z  
2023-11-23T15:58:20.818+0530 [INFO] core: post-unseal setup  
complete  
2023-11-23T15:58:20.819+0530 [INFO] core: root token  
generated  
2023-11-23T15:58:20.819+0530 [INFO] core: pre-seal teardown  
starting  
2023-11-23T15:58:20.819+0530 [INFO] rollback: stopping  
rollback manager  
2023-11-23T15:58:20.819+0530 [INFO] core: pre-seal teardown  
complete  
2023-11-23T15:58:21.116+0530 [INFO] core.cluster-  
listener.tcp: starting listener: listener_address=0.0.0.0:8201  
2023-11-23T15:58:21.116+0530 [INFO] core.cluster-listener:  
serving cluster requests: cluster_listen_address=[:]:8201  
2023-11-23T15:58:21.117+0530 [INFO] core: post-unseal setup  
starting  
2023-11-23T15:58:21.117+0530 [INFO] core: loaded wrapping  
token key  
2023-11-23T15:58:21.117+0530 [INFO] core: successfully setup  
plugin catalog: plugin-directory=""  
2023-11-23T15:58:21.119+0530 [INFO] core: successfully  
mounted: type=system version="v1.14.0+builtin.vault" path=sys/  
namespace="ID: root. Path: "  
2023-11-23T15:58:21.120+0530 [INFO] core: successfully  
mounted: type=identity version="v1.14.0+builtin.vault"  
path=identity/ namespace="ID: root. Path: "  
2023-11-23T15:58:21.120+0530 [INFO] core: successfully  
mounted: type=cubbyhole version="v1.14.0+builtin.vault"  
path=cubbyhole/ namespace="ID: root. Path: "  
2023-11-23T15:58:21.123+0530 [INFO] core: successfully  
mounted: type=token version="v1.14.0+builtin.vault"  
path=token/ namespace="ID: root. Path: "
```

```
2023-11-23T15:58:21.123+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:58:21.124+0530 [INFO] core: restoring leases
2023-11-23T15:58:21.124+0530 [INFO] identity: entities
restored
2023-11-23T15:58:21.124+0530 [INFO] identity: groups restored
2023-11-23T15:58:21.124+0530 [INFO] expiration: lease restore
complete
2023-11-23T15:58:21.125+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:58:21.125+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:58:21.125+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-11-23T15:58:21.600+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-11-23T15:58:21.690+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Upgrading to NetApp-MetroCluster-Tiebreaker-Software-1.6P1-
1.x86_64.rpm
Preparing...
##### [100%]
Updating / installing...
 1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
```

opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok

```

opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/

Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software

Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
to version 1.6P1.
Cleaning up / removing...
    2:NetApp-MetroCluster-Tiebreaker-
So##### [100%]

```

Install or upgrade to Tiebreaker 1.6

You can install Tiebreaker 1.6 or upgrade to Tiebreaker 1.6 from Tiebreaker 1.5 or 1.4.

Steps

1. Download the MetroCluster Tiebreaker 1.6 software.

[MetroCluster Tiebreaker \(Downloads\) - NetApp Support Site](#)

2. Log in to the host as the root user.
3. If you are performing an upgrade, verify the version of Tiebreaker that you are running:

The following example shows Tiebreaker 1.5.

```

[root@mcctb ~] # netapp-metrocluster-tiebreaker-software-cli
NetApp MetroCluster Tiebreaker :> version show
NetApp MetroCluster Tiebreaker 1.5: Sun Mar 13 09:59:02 IST 2022
NetApp MetroCluster Tiebreaker :> exit

```

4. Install or upgrade the Tiebreaker software.

Install Tiebreaker 1.6

Use the following steps for a new installation of Tiebreaker 1.6.

Steps

1. Run the following command at the `[root@mcctb ~] #` prompt to begin the installation:

```
sh MetroClusterTiebreakerInstall-1.6
```

The system displays the following output for a successful installation:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
Enter unix user account to use for the installation:
mcctbadminuser
Unix user account "mcctbadminuser" doesn't exist. Do you wish
to create "mcctbadminuser" user account? [Y/N]: y
useradd: warning: the home directory already exists.
Not copying any file from skel directory into it.
Creating mailbox file: File exists
Unix account "mcctbadminuser" created.
Changing password for user mcctbadminuser.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
MetroCluster Tiebreaker requires unix user account
"mcctbadminuser" to be added to the group "mcctbgrp" for admin
access.
Do you wish to add ? [Y/N]: y
Unix user account "mcctbadminuser" added to "mcctbgrp".
Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
Starting vault server...
==> Vault server configuration:

        Api Address: <api_address>
                Cgo: disabled
        Cluster Address: <cluster_address>
        Environment Variables: BASH_FUNC_which%%,
        DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
        HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
        LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
        SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,
        VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, _,
        vault_Addr, which_declare
        Go Version: go1.20.5
        Listener 1: tcp (addr: "0.0.0.0:8200", cluster
        address: "0.0.0.0:8201", max_request_duration: "1m30s",
        max_request_size: "33554432", tls: "enabled")
        Log Level:
```

```
                Mlock: supported: true, enabled: true
                Recovery Mode: false
                Storage: file
                Version: Vault v1.14.0, built 2023-06-
19T11:40:23Z
                Version Sha:
13a649f860186dffe3f3a4459814d87191efc321

==> Vault server started! Log data will stream in below:

2023-11-23T15:14:28.532+0530 [INFO] proxy environment:
http_proxy="" https_proxy="" no_proxy=""
2023-11-23T15:14:28.577+0530 [INFO] core: Initializing
version history cache for core
2023-11-23T15:14:38.552+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:14:38.552+0530 [INFO] core: seal configuration
missing, not initialized
2023-11-23T15:14:38.554+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:14:38.555+0530 [INFO] core: security barrier
initialized: stored=1 shares=5 threshold=3
2023-11-23T15:14:38.556+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:14:38.577+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:14:38.577+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:14:38.577+0530 [INFO] core: no mounts; adding
default mount table
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:14:38.578+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:14:38.581+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-11-23T15:14:38.581+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:14:38.581+0530 [INFO] core: restoring leases
2023-11-23T15:14:38.582+0530 [INFO] expiration: lease restore
```

```
complete
2023-11-23T15:14:38.582+0530 [INFO] identity: entities
restored
2023-11-23T15:14:38.582+0530 [INFO] identity: groups restored
2023-11-23T15:14:38.583+0530 [INFO] core: Recorded vault
version: vault version=1.14.0 upgrade time="2023-11-23
09:44:38.582881162 +0000 UTC" build date=2023-06-19T11:40:23Z
2023-11-23T15:14:38.583+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:14:38.998+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:14:38.999+0530 [INFO] core: root token
generated
2023-11-23T15:14:38.999+0530 [INFO] core: pre-seal teardown
starting
2023-11-23T15:14:38.999+0530 [INFO] rollback: stopping
rollback manager
2023-11-23T15:14:38.999+0530 [INFO] core: pre-seal teardown
complete
2023-11-23T15:14:39.311+0530 [INFO] core.cluster-
listener.tcp: starting listener: listener_address=0.0.0.0:8201
2023-11-23T15:14:39.311+0530 [INFO] core.cluster-listener:
serving cluster requests: cluster_listen_address=[:]:8201
2023-11-23T15:14:39.312+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:14:39.312+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:14:39.312+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:14:39.313+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:14:39.314+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-11-23T15:14:39.314+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:14:39.314+0530 [INFO] core: restoring leases
2023-11-23T15:14:39.314+0530 [INFO] identity: entities
restored
```

```

2023-11-23T15:14:39.314+0530 [INFO] expiration: lease restore
complete
2023-11-23T15:14:39.314+0530 [INFO] identity: groups restored
2023-11-23T15:14:39.315+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:14:39.316+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:14:39.316+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-11-23T15:14:39.795+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-11-23T15:14:39.885+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Installing the NetApp-MetroCluster-Tiebreaker-Software-1.6-
1.x86_64.rpm
Preparing... #
##### # [100%]

Updating / installing...

1:NetApp-MetroCluster-Tiebreaker-So#
##### # [100%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok

```

opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok

```
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/
```

```
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
Created symlink /etc/systemd/system/multi-
user.target.wants/netapp-metrocluster-tiebreaker-
software.service → /etc/systemd/system/netapp-metrocluster-
tiebreaker-software.service.
```

```
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully installed NetApp MetroCluster Tiebreaker software
version 1.6.
```

Upgrade 1.5 to 1.6

Use the following steps to upgrade the Tiebreaker 1.5 software version to Tiebreaker 1.6.

Steps

1. Run the following command at the [root@mcctb ~] # prompt to upgrade the software:

```
sh MetroClusterTiebreakerInstall-1.6
```

The system displays the following output for a successful upgrade:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok

Enter database user name : root

Please enter database password for root
Enter password:

Password updated successfully in the database.

Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
==> Vault shutdown triggered
2023-07-21T00:30:22.335+0530 [INFO] core: marked as sealed
2023-07-21T00:30:22.335+0530 [INFO] core: pre-seal teardown
starting
2023-07-21T00:30:22.335+0530 [INFO] rollback: stopping
rollback manager
2023-07-21T00:30:22.335+0530 [INFO] core: pre-seal teardown
complete
2023-07-21T00:30:22.335+0530 [INFO] core: stopping cluster
listeners
2023-07-21T00:30:22.335+0530 [INFO] core.cluster-listener:
forwarding rpc listeners stopped
2023-07-21T00:30:22.375+0530 [INFO] core.cluster-listener:
rpc listeners successfully shut down
2023-07-21T00:30:22.375+0530 [INFO] core: cluster listeners
successfully shut down
2023-07-21T00:30:22.376+0530 [INFO] core: vault is sealed
Starting vault server...
==> Vault server configuration:

    Api Address: <api_address>
        Cgo: disabled
    Cluster Address: <cluster_address>
    Environment Variables: BASH_FUNC_which%%,
    DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
    HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
    LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
```

```
SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,  
VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, _,  
vault_Addr, which_declare
```

```
Go Version: go1.20.5
```

```
Listener 1: tcp (addr: "0.0.0.0:8200", cluster  
address: "0.0.0.0:8201", max_request_duration: "1m30s",  
max_request_size: "33554432", tls: "enabled")
```

```
Log Level:
```

```
Mlock: supported: true, enabled: true
```

```
Recovery Mode: false
```

```
Storage: file
```

```
Version: Vault v1.14.0, built 2023-06-
```

```
19T11:40:23Z
```

```
Version Sha:
```

```
13a649f860186dffe3f3a4459814d87191efc321
```

```
==> Vault server started! Log data will stream in below:
```

```
2023-07-21T00:30:33.065+0530 [INFO] proxy environment:  
http_proxy="" https_proxy="" no_proxy=""  
2023-07-21T00:30:33.098+0530 [INFO] core: Initializing  
version history cache for core  
2023-07-21T00:30:43.092+0530 [INFO] core: security barrier  
not initialized  
2023-07-21T00:30:43.092+0530 [INFO] core: seal configuration  
missing, not initialized  
2023-07-21T00:30:43.094+0530 [INFO] core: security barrier  
not initialized  
2023-07-21T00:30:43.096+0530 [INFO] core: security barrier  
initialized: stored=1 shares=5 threshold=3  
2023-07-21T00:30:43.098+0530 [INFO] core: post-unseal setup  
starting  
2023-07-21T00:30:43.124+0530 [INFO] core: loaded wrapping  
token key  
2023-07-21T00:30:43.124+0530 [INFO] core: successfully setup  
plugin catalog: plugin-directory=""  
2023-07-21T00:30:43.124+0530 [INFO] core: no mounts; adding  
default mount table  
2023-07-21T00:30:43.125+0530 [INFO] core: successfully  
mounted: type=cubbyhole version="v1.14.0+builtin.vault"  
path=cubbyhole/ namespace="ID: root. Path: "  
2023-07-21T00:30:43.126+0530 [INFO] core: successfully  
mounted: type=system version="v1.14.0+builtin.vault" path=sys/  
namespace="ID: root. Path: "  
2023-07-21T00:30:43.126+0530 [INFO] core: successfully  
mounted: type=identity version="v1.14.0+builtin.vault"
```

```
path=identity/ namespace="ID: root. Path: "
2023-07-21T00:30:43.129+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-07-21T00:30:43.130+0530 [INFO] rollback: starting
rollback manager
2023-07-21T00:30:43.130+0530 [INFO] core: restoring leases
2023-07-21T00:30:43.130+0530 [INFO] identity: entities
restored
2023-07-21T00:30:43.130+0530 [INFO] identity: groups restored
2023-07-21T00:30:43.131+0530 [INFO] core: usage gauge
collection is disabled
2023-07-21T00:30:43.131+0530 [INFO] expiration: lease restore
complete
2023-07-21T00:30:43.131+0530 [INFO] core: Recorded vault
version: vault version=1.14.0 upgrade time="2023-07-20
19:00:43.131158543 +0000 UTC" build date=2023-06-19T11:40:23Z
2023-07-21T00:30:43.371+0530 [INFO] core: post-unseal setup
complete
2023-07-21T00:30:43.371+0530 [INFO] core: root token
generated
2023-07-21T00:30:43.371+0530 [INFO] core: pre-seal teardown
starting
2023-07-21T00:30:43.371+0530 [INFO] rollback: stopping
rollback manager
2023-07-21T00:30:43.372+0530 [INFO] core: pre-seal teardown
complete
2023-07-21T00:30:43.694+0530 [INFO] core.cluster-
listener.tcp: starting listener: listener_address=0.0.0.0:8201
2023-07-21T00:30:43.695+0530 [INFO] core.cluster-listener:
serving cluster requests: cluster_listen_address=[:]:8201
2023-07-21T00:30:43.695+0530 [INFO] core: post-unseal setup
starting
2023-07-21T00:30:43.696+0530 [INFO] core: loaded wrapping
token key
2023-07-21T00:30:43.696+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-07-21T00:30:43.697+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-07-21T00:30:43.698+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-07-21T00:30:43.698+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
```

```
2023-07-21T00:30:43.701+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
path=token/ namespace="ID: root. Path: "
2023-07-21T00:30:43.701+0530 [INFO] rollback: starting
rollback manager
2023-07-21T00:30:43.702+0530 [INFO] core: restoring leases
2023-07-21T00:30:43.702+0530 [INFO] identity: entities
restored
2023-07-21T00:30:43.702+0530 [INFO] expiration: lease restore
complete
2023-07-21T00:30:43.702+0530 [INFO] identity: groups restored
2023-07-21T00:30:43.702+0530 [INFO] core: usage gauge
collection is disabled
2023-07-21T00:30:43.703+0530 [INFO] core: post-unseal setup
complete
2023-07-21T00:30:43.703+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-07-21T00:30:44.226+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-07-21T00:30:44.315+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Upgrading to NetApp-MetroCluster-Tiebreaker-Software-1.6-
1.x86_64.rpm
Preparing...
##### [100%]
Updating / installing...
 1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
```

opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok

```
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok
opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/
```

```
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
```

```
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
to version 1.6.
Cleaning up / removing...
  2:NetApp-MetroCluster-Tiebreaker-
So##### [100%]
```

Upgrade 1.4 to 1.6

Use the following steps to upgrade the Tiebreaker 1.4 software version to Tiebreaker 1.6.

Steps

1. Run the following command at the [root@mcctb ~] # prompt to upgrade the software:

```
sh MetroClusterTiebreakerInstall-1.6
```

The system displays the following output for a successful upgrade:

Example

```
Extracting the MetroCluster Tiebreaker installation/upgrade
archive
Install digest hash is Ok
Performing the MetroCluster Tiebreaker code signature check
Install code signature is Ok
Enter unix user account to use for the installation:
mcctbuseradmin1
Unix user account "mcctbuseradmin1" doesn't exist. Do you wish
to create "mcctbuseradmin1" user account? [Y/N]: y
Unix account "mcctbuseradmin1" created.
Changing password for user mcctbuseradmin1.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.

Enter database user name : root

Please enter database password for root
Enter password:

Password updated successfully in the database.

MetroCluster Tiebreaker requires unix user account
"mcctbuseradmin1" to be added to the group "mcctbgrp" for
admin access.
Do you wish to add ? [Y/N]: y
Unix user account "mcctbuseradmin1" added to "mcctbgrp".
Do you wish to generate your own public-private key pair for
encrypting audit log? [Y/N]: y
Generating public-private key pair...
Configuring Vault...
Starting vault server...
==> Vault server configuration:

        Api Address: <api_address>
                Cgo: disabled
        Cluster Address: <cluster_address>
        Environment Variables: BASH_FUNC_which%,
        DBUS_SESSION_BUS_ADDRESS, GODEBUG, HISTCONTROL, HISTSIZE,
        HOME, HOSTNAME, HOST_ACCOUNT, LANG, LESSOPEN, LOGNAME,
        LS_COLORS, MAIL, PATH, PWD, SHELL, SHLVL, SSH_CLIENT,
        SSH_CONNECTION, SSH_TTY, STAF_TEMP_DIR, TERM, USER,
        VAULT_ADDR, VAULT_TOKEN, XDG_RUNTIME_DIR, XDG_SESSION_ID, __,
        vault_Addr, which_declare
```

```
Go Version: go1.20.5
Listener 1: tcp (addr: "0.0.0.0:8200", cluster
address: "0.0.0.0:8201", max_request_duration: "1m30s",
max_request_size: "33554432", tls: "enabled")
Log Level:
Mlock: supported: true, enabled: true
Recovery Mode: false
Storage: file
Version: Vault v1.14.0, built 2023-06-
19T11:40:23Z
Version Sha:
13a649f860186dffe3f3a4459814d87191efc321
```

==> Vault server started! Log data will stream in below:

```
2023-11-23T15:58:10.400+0530 [INFO] proxy environment:
http_proxy="" https_proxy="" no_proxy=""
2023-11-23T15:58:10.432+0530 [INFO] core: Initializing
version history cache for core
2023-11-23T15:58:20.422+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:58:20.422+0530 [INFO] core: seal configuration
missing, not initialized
2023-11-23T15:58:20.424+0530 [INFO] core: security barrier
not initialized
2023-11-23T15:58:20.425+0530 [INFO] core: security barrier
initialized: stored=1 shares=5 threshold=3
2023-11-23T15:58:20.427+0530 [INFO] core: post-unseal setup
starting
2023-11-23T15:58:20.448+0530 [INFO] core: loaded wrapping
token key
2023-11-23T15:58:20.448+0530 [INFO] core: successfully setup
plugin catalog: plugin-directory=""
2023-11-23T15:58:20.448+0530 [INFO] core: no mounts; adding
default mount table
2023-11-23T15:58:20.449+0530 [INFO] core: successfully
mounted: type=cubbyhole version="v1.14.0+builtin.vault"
path=cubbyhole/ namespace="ID: root. Path: "
2023-11-23T15:58:20.449+0530 [INFO] core: successfully
mounted: type=system version="v1.14.0+builtin.vault" path=sys/
namespace="ID: root. Path: "
2023-11-23T15:58:20.449+0530 [INFO] core: successfully
mounted: type=identity version="v1.14.0+builtin.vault"
path=identity/ namespace="ID: root. Path: "
2023-11-23T15:58:20.451+0530 [INFO] core: successfully
mounted: type=token version="v1.14.0+builtin.vault"
```

```
path=token/ namespace="ID: root. Path: "  
2023-11-23T15:58:20.452+0530 [INFO] rollback: starting  
rollback manager  
2023-11-23T15:58:20.452+0530 [INFO] core: restoring leases  
2023-11-23T15:58:20.453+0530 [INFO] identity: entities  
restored  
2023-11-23T15:58:20.453+0530 [INFO] identity: groups restored  
2023-11-23T15:58:20.453+0530 [INFO] expiration: lease restore  
complete  
2023-11-23T15:58:20.453+0530 [INFO] core: usage gauge  
collection is disabled  
2023-11-23T15:58:20.453+0530 [INFO] core: Recorded vault  
version: vault version=1.14.0 upgrade time="2023-11-23  
10:28:20.453481904 +0000 UTC" build date=2023-06-19T11:40:23Z  
2023-11-23T15:58:20.818+0530 [INFO] core: post-unseal setup  
complete  
2023-11-23T15:58:20.819+0530 [INFO] core: root token  
generated  
2023-11-23T15:58:20.819+0530 [INFO] core: pre-seal teardown  
starting  
2023-11-23T15:58:20.819+0530 [INFO] rollback: stopping  
rollback manager  
2023-11-23T15:58:20.819+0530 [INFO] core: pre-seal teardown  
complete  
2023-11-23T15:58:21.116+0530 [INFO] core.cluster-  
listener.tcp: starting listener: listener_address=0.0.0.0:8201  
2023-11-23T15:58:21.116+0530 [INFO] core.cluster-listener:  
serving cluster requests: cluster_listen_address=[:]:8201  
2023-11-23T15:58:21.117+0530 [INFO] core: post-unseal setup  
starting  
2023-11-23T15:58:21.117+0530 [INFO] core: loaded wrapping  
token key  
2023-11-23T15:58:21.117+0530 [INFO] core: successfully setup  
plugin catalog: plugin-directory=""  
2023-11-23T15:58:21.119+0530 [INFO] core: successfully  
mounted: type=system version="v1.14.0+builtin.vault" path=sys/  
namespace="ID: root. Path: "  
2023-11-23T15:58:21.120+0530 [INFO] core: successfully  
mounted: type=identity version="v1.14.0+builtin.vault"  
path=identity/ namespace="ID: root. Path: "  
2023-11-23T15:58:21.120+0530 [INFO] core: successfully  
mounted: type=cubbyhole version="v1.14.0+builtin.vault"  
path=cubbyhole/ namespace="ID: root. Path: "  
2023-11-23T15:58:21.123+0530 [INFO] core: successfully  
mounted: type=token version="v1.14.0+builtin.vault"  
path=token/ namespace="ID: root. Path: "
```

```
2023-11-23T15:58:21.123+0530 [INFO] rollback: starting
rollback manager
2023-11-23T15:58:21.124+0530 [INFO] core: restoring leases
2023-11-23T15:58:21.124+0530 [INFO] identity: entities
restored
2023-11-23T15:58:21.124+0530 [INFO] identity: groups restored
2023-11-23T15:58:21.124+0530 [INFO] expiration: lease restore
complete
2023-11-23T15:58:21.125+0530 [INFO] core: usage gauge
collection is disabled
2023-11-23T15:58:21.125+0530 [INFO] core: post-unseal setup
complete
2023-11-23T15:58:21.125+0530 [INFO] core: vault is unsealed
Success! Uploaded policy: mcctb-policy
2023-11-23T15:58:21.600+0530 [INFO] core: enabled credential
backend: path=appprole/ type=appprole version=""
Success! Enabled approle auth method at: approle/
2023-11-23T15:58:21.690+0530 [INFO] core: successful mount:
namespace="" path=mcctb/ type=kv version=""
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/appprole/role/mcctb-app
Upgrading to NetApp-MetroCluster-Tiebreaker-Software-1.6-
1.x86_64.rpm
Preparing...
##### [100%]
Updating / installing...
 1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Performing file integrity check
etc/cron.weekly/metrocluster-tiebreaker-support is Ok
etc/cron.weekly/metrocluster-tiebreaker-support-cov is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software is Ok
etc/init.d/netapp-metrocluster-tiebreaker-software-cov is Ok
etc/logrotate.d/mcctb is Ok
opt/netapp/mcctb/lib/common/activation-1.1.1.jar is Ok
opt/netapp/mcctb/lib/common/aopalliance.jar is Ok
opt/netapp/mcctb/lib/common/args4j.jar is Ok
opt/netapp/mcctb/lib/common/aspectjrt.jar is Ok
opt/netapp/mcctb/lib/common/aspectjweaver.jar is Ok
opt/netapp/mcctb/lib/common/asup.jar is Ok
opt/netapp/mcctb/lib/common/bcpkix-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk15on.jar is Ok
opt/netapp/mcctb/lib/common/bcprov-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bctls-fips-1.0.13.jar is Ok
opt/netapp/mcctb/lib/common/bctls-jdk18on.jar is Ok
opt/netapp/mcctb/lib/common/bcutil-jdk18on.jar is Ok
```

opt/netapp/mcctb/lib/common/cglib.jar is Ok
opt/netapp/mcctb/lib/common/commons-codec.jar is Ok
opt/netapp/mcctb/lib/common/commons-collections4.jar is Ok
opt/netapp/mcctb/lib/common/commons-compress.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.jar is Ok
opt/netapp/mcctb/lib/common/commons-daemon.src.jar is Ok
opt/netapp/mcctb/lib/common/commons-dbcp2.jar is Ok
opt/netapp/mcctb/lib/common/commons-io.jar is Ok
opt/netapp/mcctb/lib/common/commons-lang3.jar is Ok
opt/netapp/mcctb/lib/common/commons-logging.jar is Ok
opt/netapp/mcctb/lib/common/commons-pool2.jar is Ok
opt/netapp/mcctb/lib/common/guava.jar is Ok
opt/netapp/mcctb/lib/common/httpclient.jar is Ok
opt/netapp/mcctb/lib/common/httpcore.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.activation.jar is Ok
opt/netapp/mcctb/lib/common/jakarta.xml.bind-api.jar is Ok
opt/netapp/mcctb/lib/common/java-xmlbuilder.jar is Ok
opt/netapp/mcctb/lib/common/javax.inject.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-api-2.3.1.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-core.jar is Ok
opt/netapp/mcctb/lib/common/jaxb-impl.jar is Ok
opt/netapp/mcctb/lib/common/jline.jar is Ok
opt/netapp/mcctb/lib/common/jna.jar is Ok
opt/netapp/mcctb/lib/common/joda-time.jar is Ok
opt/netapp/mcctb/lib/common/jsch.jar is Ok
opt/netapp/mcctb/lib/common/json.jar is Ok
opt/netapp/mcctb/lib/common/jsvc.zip is Ok
opt/netapp/mcctb/lib/common/junixsocket-common.jar is Ok
opt/netapp/mcctb/lib/common/junixsocket-native-common.jar is Ok
Ok
opt/netapp/mcctb/lib/common/logback-classic.jar is Ok
opt/netapp/mcctb/lib/common/logback-core.jar is Ok
opt/netapp/mcctb/lib/common/mail-1.6.2.jar is Ok
opt/netapp/mcctb/lib/common/mariadb-java-client.jar is Ok
opt/netapp/mcctb/lib/common/mcctb-mib.jar is Ok
opt/netapp/mcctb/lib/common/mcctb.jar is Ok
opt/netapp/mcctb/lib/common/mockito-core.jar is Ok
opt/netapp/mcctb/lib/common/slf4j-api.jar is Ok
opt/netapp/mcctb/lib/common/snmp4j.jar is Ok
opt/netapp/mcctb/lib/common/spring-aop.jar is Ok
opt/netapp/mcctb/lib/common/spring-beans.jar is Ok
opt/netapp/mcctb/lib/common/spring-context-support.jar is Ok
opt/netapp/mcctb/lib/common/spring-context.jar is Ok
opt/netapp/mcctb/lib/common/spring-core.jar is Ok
opt/netapp/mcctb/lib/common/spring-expression.jar is Ok
opt/netapp/mcctb/lib/common/spring-web.jar is Ok

```

opt/netapp/mcctb/lib/common/vault-java-driver.jar is Ok
opt/netapp/mcctb/lib/common/xz.jar is Ok
opt/netapp/mcctb/lib/org.jacoco.agent-0.8.8-runtime.jar is Ok
opt/netapp/mcctb/bin/mcctb-asup-invoke is Ok
opt/netapp/mcctb/bin/mcctb_postrotate is Ok
opt/netapp/mcctb/bin/netapp-metrocluster-tiebreaker-software-
cli is Ok
/

Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software

Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
to version 1.6.
Cleaning up / removing...
    2:NetApp-MetroCluster-Tiebreaker-
So##### [100%]

```

Install Tiebreaker 1.5

Configure admin access to ONTAP API and SSH

You can configure admin access to ONTAP API and SSH.

Steps

1. Create an admin user with ONTAP API access: `security login create -user-or-group-name mcctb -application ontapi -authentication-method password`
2. Create an admin user with SSH access: `security login create -user-or-group-name mcctb -application ssh -authentication-method password`
3. Verify that the new admin users are created: `security login show`
4. Repeat these steps on the partner cluster.



Administrator authentication and RBAC is implemented.

Install MetroCluster Tiebreaker 1.5 dependencies

Related information

Depending on your host Linux operating system, you must install a MySQL or MariaDB server before installing or upgrading the Tiebreaker software.

Steps

1. [Install JDK](#)
2. [Install and configure Vault](#)
3. Install MySQL or MariaDB server:

If the Linux host is	Then...
Red Hat Enterprise Linux 7/CentOS 7	Install MySQL Server 5.5.30 or later and 5.6.x versions on Red Hat Enterprise Linux 7 or CentOS 7
Red Hat Enterprise Linux 8	Install MariaDB server on Red Hat Enterprise Linux 8

Install JDK

You must install JDK on your host system before installing or upgrading the Tiebreaker software. Tiebreaker 1.5 and later supports OpenJDK 17, 18, or 19.

Steps

1. Log in as a "root" user or a sudo user that can change to advanced privilege mode.

```
login as: root
root@mcctb's password:
Last login: Fri Jan  8 21:33:00 2017 from host.domain.com
```

2. Check for available JDK versions:

```
yum search openjdk
```

3. Install JDK 17,18, or 19.

The following command installs JDK 17:

```
yum install java-17-openjdk
```

4. Verify the installation:

```
java -version
```

A successful installation displays the following output:

```
openjdk version "17.0.2" 2022-01-18 LTS
OpenJDK Runtime Environment 21.9 (build 17.0.2+8-LTS)
OpenJDK 64-Bit Server VM 21.9 (build 17.0.2+8-LTS, mixed mode, sharing)
```

Install and configure Vault

If you do not have or want to use the local Vault server, you must install Vault.

You can refer to this standard procedure for installing Vault, or refer to the Hashicorp installation instructions for alternative guidelines.



If you have a Vault server in your network, you can configure the MetroCluster Tiebreaker host to use that Vault installation. If you do this, you do not need to install Vault on the host.

Steps

1. Navigate to the `/bin` directory:

```
[root@mcctb] cd /bin
```

2. Download the Vault zip file.

```
[root@mcctb /bin]# curl -sO  
https://releases.hashicorp.com/vault/1.12.2/vault_1.12.2_linux_amd64.zip
```

3. Unzip the Vault file.

```
[root@mcctb /bin]# unzip vault_1.12.2_linux_amd64.zip  
Archive:  vault_1.12.2_linux_amd64.zip  
  inflating: vault
```

4. Verify the installation.

```
[root@mcctb /bin]# vault -version  
Vault v1.12.2 (415e1fe3118eebd5df6cb60d13defdc01aa17b03), built 2022-11-  
23T12:53:46Z
```

5. Navigate to the `/root` directory:

```
[root@mcctb /bin] cd /root
```

6. Create a Vault configuration file under the `/root` directory.

At the `[root@mcctb ~]` prompt, copy and run the following command to create the `config.hcl` file:

```
# cat > config.hcl << EOF
storage "file" {
  address = "127.0.0.1:8500"
  path    = "/mcctb_vdata/data"
}
listener "tcp" {
  address      = "127.0.0.1:8200"
  tls_disable = 1
}
EOF
```

7. Start the Vault server:

```
[root@mcctb ~] vault server -config config.hcl &
```

8. Export the Vault address.

```
[root@mcctb ~]# export VAULT_ADDR="http://127.0.0.1:8200"
```

9. Initialize Vault.

```
[root@mcctb ~]# vault operator init
2022-12-15T14:57:22.113+0530 [INFO] core: security barrier not
initialized
2022-12-15T14:57:22.113+0530 [INFO] core: seal configuration missing,
not initialized
2022-12-15T14:57:22.114+0530 [INFO] core: security barrier not
initialized
2022-12-15T14:57:22.116+0530 [INFO] core: security barrier initialized:
stored=1 shares=5 threshold=3
2022-12-15T14:57:22.118+0530 [INFO] core: post-unseal setup starting
2022-12-15T14:57:22.137+0530 [INFO] core: loaded wrapping token key
2022-12-15T14:57:22.137+0530 [INFO] core: Recorded vault version: vault
version=1.12.2 upgrade time="2022-12-15 09:27:22.137200412 +0000 UTC"
build date=2022-11-23T12:53:46Z
2022-12-15T14:57:22.137+0530 [INFO] core: successfully setup plugin
catalog: plugin-directory=""
2022-12-15T14:57:22.137+0530 [INFO] core: no mounts; adding default
mount table
2022-12-15T14:57:22.143+0530 [INFO] core: successfully mounted backend:
type=cubbyhole version="" path=cubbyhole/
2022-12-15T14:57:22.144+0530 [INFO] core: successfully mounted backend:
type=system version="" path=sys/
```

```
2022-12-15T14:57:22.144+0530 [INFO] core: successfully mounted backend:
type=identity version="" path=identity/
2022-12-15T14:57:22.148+0530 [INFO] core: successfully enabled
credential backend: type=token version="" path=token/ namespace="ID:
root. Path: "
2022-12-15T14:57:22.149+0530 [INFO] rollback: starting rollback manager
2022-12-15T14:57:22.149+0530 [INFO] core: restoring leases
2022-12-15T14:57:22.150+0530 [INFO] expiration: lease restore complete
2022-12-15T14:57:22.150+0530 [INFO] identity: entities restored
2022-12-15T14:57:22.150+0530 [INFO] identity: groups restored
2022-12-15T14:57:22.151+0530 [INFO] core: usage gauge collection is
disabled
2022-12-15T14:57:23.385+0530 [INFO] core: post-unseal setup complete
2022-12-15T14:57:23.387+0530 [INFO] core: root token generated
2022-12-15T14:57:23.387+0530 [INFO] core: pre-seal teardown starting
2022-12-15T14:57:23.387+0530 [INFO] rollback: stopping rollback manager
2022-12-15T14:57:23.387+0530 [INFO] core: pre-seal teardown complete
Unseal Key 1: <unseal_key_1_id>
Unseal Key 2: <unseal_key_2_id>
Unseal Key 3: <unseal_key_3_id>
Unseal Key 4: <unseal_key_4_id>
Unseal Key 5: <unseal_key_5_id>

Initial Root Token: <initial_root_token_id>
```

Vault initialized with 5 key shares and a key threshold of 3. Please securely distribute the key shares printed above. When the Vault is re-sealed, restarted, or stopped, you must supply at least 3 of these keys to unseal it before it can start servicing requests.

Vault does not store the generated root key. Without at least 3 keys to reconstruct the root key, Vault will remain permanently sealed!

It is possible to generate new unseal keys, provided you have a quorum of existing unseal keys shares. See "vault operator rekey" for more information.



You must record and store the key IDs and initial root token in a secure location for use later in the procedure.

10. Export the Vault root token.

```
[root@mcctb ~]# export VAULT_TOKEN="<initial_root_token_id>"
```

11. Unseal Vault by using any three of the five keys that were created.

You must run the `vault operator unseal` command for each of the three keys:

a. Unseal vault by using the first key:

```
[root@mcctb ~]# vault operator unseal
Unseal Key (will be hidden):
Key          Value
---          -
Seal Type    shamir
Initialized   true
Sealed       true
Total Shares  5
Threshold    3
Unseal Progress 1/3
Unseal Nonce <unseal_key_1_id>
Version      1.12.2
Build Date   2022-11-23T12:53:46Z
Storage Type file
HA Enabled   false
```

b. Unseal vault by using the second key:

```
[root@mcctb ~]# vault operator unseal
Unseal Key (will be hidden):
Key          Value
---          -
Seal Type    shamir
Initialized   true
Sealed       true
Total Shares  5
Threshold    3
Unseal Progress 2/3
Unseal Nonce <unseal_key_2_id>
Version      1.12.2
Build Date   2022-11-23T12:53:46Z
Storage Type file
HA Enabled   false
```

c. Unseal vault by using the third key:

```

[root@mcctb ~]# vault operator unseal
Unseal Key (will be hidden):
2022-12-15T15:15:00.980+0530 [INFO] core.cluster-listener.tcp:
starting listener: listener_address=127.0.0.1:8201
2022-12-15T15:15:00.980+0530 [INFO] core.cluster-listener: serving
cluster requests: cluster_listen_address=127.0.0.1:8201
2022-12-15T15:15:00.981+0530 [INFO] core: post-unseal setup starting
2022-12-15T15:15:00.981+0530 [INFO] core: loaded wrapping token key
2022-12-15T15:15:00.982+0530 [INFO] core: successfully setup plugin
catalog: plugin-directory=""
2022-12-15T15:15:00.983+0530 [INFO] core: successfully mounted
backend: type=system version="" path=sys/
2022-12-15T15:15:00.984+0530 [INFO] core: successfully mounted
backend: type=identity version="" path=identity/
2022-12-15T15:15:00.984+0530 [INFO] core: successfully mounted
backend: type=cubbyhole version="" path=cubbyhole/
2022-12-15T15:15:00.986+0530 [INFO] core: successfully enabled
credential backend: type=token version="" path=token/ namespace="ID:
root. Path: "
2022-12-15T15:15:00.986+0530 [INFO] rollback: starting rollback
manager
2022-12-15T15:15:00.987+0530 [INFO] core: restoring leases
2022-12-15T15:15:00.987+0530 [INFO] expiration: lease restore
complete
2022-12-15T15:15:00.987+0530 [INFO] identity: entities restored
2022-12-15T15:15:00.987+0530 [INFO] identity: groups restored
2022-12-15T15:15:00.988+0530 [INFO] core: usage gauge collection is
disabled
2022-12-15T15:15:00.989+0530 [INFO] core: post-unseal setup complete
2022-12-15T15:15:00.989+0530 [INFO] core: vault is unsealed
Key          Value
---          -
Seal Type    shamir
Initialized   true
Sealed       false
Total Shares  5
Threshold    3
Version      1.12.2
Build Date   2022-11-23T12:53:46Z
Storage Type  file
Cluster Name  vault-cluster
Cluster ID    <cluster_id>
HA Enabled    false

```

12. Verify that the Vault sealed status is false.

```
[root@mcctb ~]# vault status
Key          Value
---          -
Seal Type    shamir
Initialized  true
Sealed       false
Total Shares 5
Threshold    3
Version      1.12.2
Build Date   2022-11-23T12:53:46Z
Storage Type file
Cluster Name vault-cluster
Cluster ID   <cluster_id>
HA Enabled   false
```

13. Configure the Vault service to start on boot.

- a. Run the following command: `cd /etc/systemd/system`

```
[root@mcctb ~]# cd /etc/systemd/system
```

- b. At the `[root@mcctb system]` prompt, copy and run the following command to create the Vault service file.

```
# cat > vault.service << EOF
[Unit]
Description=Vault Service
After=mariadb.service

[Service]
Type=forking
ExecStart=/usr/bin/vault server -config /root/config.hcl &
Restart=on-failure

[Install]
WantedBy=multi-user.target
EOF
```

- c. Run the following command: `systemctl daemon-reload`

```
[root@mcctb system]# systemctl daemon-reload
```

- d. Run the following command: `systemctl enable vault.service`

```
[root@mcctb system]# systemctl enable vault.service
Created symlink /etc/systemd/system/multi-
user.target.wants/vault.service → /etc/systemd/system/vault.service.
```



You are prompted to use this feature during the installation of MetroCluster Tiebreaker. If you want to change the method to unseal Vault, then you need to uninstall and reinstall the MetroCluster Tiebreaker software.

Install MySQL Server 5.5.30 or later and 5.6.x versions on Red Hat Enterprise Linux 7 or CentOS 7

You must install MySQL Server 5.5.30 or later and 5.6.x version on your host system before installing or upgrading the Tiebreaker software. For Red Hat Enterprise Linux 8, [Install MariaDB server](#).

Steps

1. Log in as a root user or a sudo user that can change to advanced privilege mode.

```
login as: root
root@mcctb's password:
Last login: Fri Jan  8 21:33:00 2016 from host.domain.com
```

2. Add the MySQL repository to your host system:

```
[root@mcctb ~]# yum localinstall https://dev.mysql.com/get/mysql57-community-
release-el6-11.noarch.rpm
```

```

Loaded plugins: product-id, refresh-packagekit, security, subscription-
manager
Setting up Local Package Process
Examining /var/tmp/yum-root-LLUw0r/mysql-community-release-el6-
5.noarch.rpm: mysql-community-release-el6-5.noarch
Marking /var/tmp/yum-root-LLUw0r/mysql-community-release-el6-
5.noarch.rpm to be installed
Resolving Dependencies
--> Running transaction check
---> Package mysql-community-release.noarch 0:el6-5 will be installed
--> Finished Dependency Resolution
Dependencies Resolved

=====
=====
Package                Arch    Version
                        Repository

Size
=====
=====
Installing:
mysql-community-release
                        noarch el6-5 /mysql-community-release-el6-
5.noarch 4.3 k
Transaction Summary
=====
=====
Install                1 Package(s)
Total size: 4.3 k
Installed size: 4.3 k
Is this ok [y/N]: y
Downloading Packages:
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : mysql-community-release-el6-5.noarch
1/1
  Verifying  : mysql-community-release-el6-5.noarch
1/1
Installed:
  mysql-community-release.noarch 0:el6-5
Complete!

```

3. Disable the MySQL 57 repository:

```
[root@mcctb ~]# yum-config-manager --disable mysql57-community
```

4. Enable the MySQL 56 repository:

```
[root@mcctb ~]# yum-config-manager --enable mysql56-community
```

5. Enable the repository:

```
[root@mcctb ~]# yum repolist enabled | grep "mysql.-community."
```

```
mysql-connectors-community           MySQL Connectors Community
21
mysql-tools-community                MySQL Tools Community
35
mysql56-community                    MySQL 5.6 Community Server
231
```

6. Install the MySQL Community server:

```
[root@mcctb ~]# yum install mysql-community-server
```

```
Loaded plugins: product-id, refresh-packagekit, security, subscription-
manager
This system is not registered to Red Hat Subscription Management. You
can use subscription-manager
to register.
Setting up Install Process
Resolving Dependencies
--> Running transaction check
.....Output truncated.....
---> Package mysql-community-libs-compat.x86_64 0:5.6.29-2.el6 will be
obsoleting
--> Finished Dependency Resolution
Dependencies Resolved

=====
=====
Package                               Arch    Version           Repository
Size
=====
Installing:
mysql-community-client                x86_64  5.6.29-2.el6     mysql56-community
18 M
    replacing mysql.x86_64 5.1.71-1.el6
mysql-community-libs                   x86_64  5.6.29-2.el6     mysql56-community
1.9 M
```

```
replacing mysql-libs.x86_64 5.1.71-1.el6
mysql-community-libs-compat x86_64 5.6.29-2.el6 mysql56-community
1.6 M
replacing mysql-libs.x86_64 5.1.71-1.el6
mysql-community-server x86_64 5.6.29-2.el6 mysql56-community
53 M
replacing mysql-server.x86_64 5.1.71-1.el6
Installing for dependencies:
mysql-community-common x86_64 5.6.29-2.el6 mysql56-community
308 k
```

Transaction Summary

```
=====
=====
```

```
Install          5 Package(s)
Total download size: 74 M
```

Is this ok [y/N]: y

Downloading Packages:

```
(1/5): mysql-community-client-5.6.29-2.el6.x86_64.rpm      | 18 MB
00:28
(2/5): mysql-community-common-5.6.29-2.el6.x86_64.rpm     | 308 kB
00:01
(3/5): mysql-community-libs-5.6.29-2.el6.x86_64.rpm      | 1.9 MB
00:05
(4/5): mysql-community-libs-compat-5.6.29-2.el6.x86_64.rpm | 1.6 MB
00:05
(5/5): mysql-community-server-5.6.29-2.el6.x86_64.rpm    | 53 MB
03:42
```

```
-----
-----
```

```
Total                                     289 kB/s | 74 MB
04:24
```

warning: rpmts_HdrFromFdno: Header V3 DSA/SHA1 Signature, key ID <key_id> NOKEY

Retrieving key from file:/etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

Importing GPG key 0x5072E1F5:

Userid : MySQL Release Engineering <mysql-build@oss.oracle.com>

Package: mysql-community-release-el6-5.noarch

(@/mysql-community-release-el6-5.noarch)

From : file:/etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

Is this ok [y/N]: y

Running rpm_check_debug

Running Transaction Test

Transaction Test Succeeded

Running Transaction

Installing : mysql-community-common-5.6.29-2.el6.x86_64

```
....Output truncated....
```

```
1.el6.x86_64
```

```
7/8
```

```
Verifying : mysql-5.1.71-1.el6.x86_64
```

```
8/8
```

```
Installed:
```

```
mysql-community-client.x86_64 0:5.6.29-2.el6
```

```
mysql-community-libs.x86_64 0:5.6.29-2.el6
```

```
mysql-community-libs-compat.x86_64 0:5.6.29-2.el6
```

```
mysql-community-server.x86_64 0:5.6.29-2.el6
```

```
Dependency Installed:
```

```
mysql-community-common.x86_64 0:5.6.29-2.el6
```

```
Replaced:
```

```
mysql.x86_64 0:5.1.71-1.el6 mysql-libs.x86_64 0:5.1.71-1.el6
```

```
mysql-server.x86_64 0:5.1.71-1.el6
```

```
Complete!
```

7. Start MySQL server:

```
[root@mcctb ~]# service mysqld start
```

```
Initializing MySQL database: 2016-04-05 19:44:38 0 [Warning] TIMESTAMP
with implicit DEFAULT value is deprecated. Please use
--explicit_defaults_for_timestamp server option (see documentation
for more details).
2016-04-05 19:44:38 0 [Note] /usr/sbin/mysqld (mysqld 5.6.29)
starting as process 2487 ...
2016-04-05 19:44:38 2487 [Note] InnoDB: Using atomics to ref count
buffer pool pages
2016-04-05 19:44:38 2487 [Note] InnoDB: The InnoDB memory heap is
disabled
....Output truncated....
2016-04-05 19:44:42 2509 [Note] InnoDB: Shutdown completed; log sequence
number 1625987
```

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER!
To do so, start the server, then issue the following commands:

```
/usr/bin/mysqladmin -u root password 'new-password'
/usr/bin/mysqladmin -u root -h mcctb password 'new-password'
```

Alternatively, you can run:

```
/usr/bin/mysql_secure_installation
```

which will also give you the option of removing the test
databases and anonymous user created by default. This is
strongly recommended for production servers.

.....Output truncated.....

```
WARNING: Default config file /etc/my.cnf exists on the system
This file will be read by default by the MySQL server
If you do not want to use this, either remove it, or use the
--defaults-file argument to mysqld_safe when starting the server
```

```
Starting mysqld: [ OK ]
```

8. Confirm that MySQL server is running:

```
[root@mcctb ~]# service mysqld status
```

```
mysqld (pid 2739) is running...
```

9. Configure security and password settings:

```
[root@mcctb ~]# mysql_secure_installation
```

NOTE: RUNNING ALL PARTS OF THIS SCRIPT IS RECOMMENDED FOR ALL MySQL
SERVERS IN PRODUCTION USE! PLEASE READ EACH STEP CAREFULLY!

In order to log into MySQL to secure it, we'll need the current password for the root user. If you've just installed MySQL, and you haven't set the root password yet, the password will be blank, so you should just press enter here.

Enter current password for root (enter for none): <== on default
install

hit enter here

OK, successfully used password, moving on...

Setting the root password ensures that nobody can log into the MySQL root user without the proper authorization.

Set root password? [Y/n] y

New password:

Re-enter new password:

Password updated successfully!

Reloading privilege tables..

... Success!

By default, a MySQL installation has an anonymous user, allowing anyone to log into MySQL without having to have a user account created for them. This is intended only for testing, and to make the installation go a bit smoother. You should remove them before moving into a production environment.

Remove anonymous users? [Y/n] y

... Success!

Normally, root should only be allowed to connect from 'localhost'. This ensures that someone cannot guess at the root password from the network.

Disallow root login remotely? [Y/n] y

... Success!

By default, MySQL comes with a database named 'test' that anyone can access. This is also intended only for testing, and should be removed before moving into a production environment.

Remove test database and access to it? [Y/n] y

- Dropping test database...

ERROR 1008 (HY000) at line 1: Can't drop database 'test';

```
database doesn't exist
```

```
... Failed! Not critical, keep moving...  
- Removing privileges on test database...  
... Success!
```

Reloading the privilege tables will ensure that all changes made so far will take effect immediately.

```
Reload privilege tables now? [Y/n] y
```

```
... Success!
```

All done! If you've completed all of the above steps, your MySQL installation should now be secure.

Thanks for using MySQL!

Cleaning up...

10. Verify that the MySQL login is working:

```
[root@mcctb ~]# mysql -u root -p
```

```
Enter password: <configured_password>
```

```
Welcome to the MySQL monitor. Commands end with ; or \g.
```

```
Your MySQL connection id is 17
```

```
Server version: 5.6.29 MySQL Community Server (GPL)
```

```
Copyright (c) 2000, 2016, Oracle and/or its affiliates. All rights reserved.
```

```
Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.
```

```
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
```

```
mysql>
```

If the MySQL login is working, the output will end at the `mysql>` prompt.

Enable the MySQL autostart setting

You should verify that the autostart feature is turned on for the MySQL daemon. Turning on the MySQL daemon automatically restarts MySQL if the system on which the MetroCluster Tiebreaker software resides reboots. If the MySQL daemon is not running, the Tiebreaker software continues running, but it cannot be restarted and configuration changes cannot be made.

Step

1. Verify that MySQL is enabled to autostart when booted:

```
[root@mcctb ~]# systemctl list-unit-files mysqld.service
```

```
UNIT FILE           State
-----
mysqld.service     enabled
```

If MySQL is not enabled to autostart when booted, see the MySQL documentation to enable the autostart feature for your installation.

Install MariaDB server on Red Hat Enterprise Linux 8

You must install MariaDB server on your host system before installing or upgrading the Tiebreaker software. For Red Hat Enterprise Linux 7 or CentOS 7, [Install MySQL Server](#).

Before you begin

Your host system must be running on Red Hat Enterprise Linux (RHEL) 8.

Steps

1. Log in as a `root` user or a user that can `sudo` to advanced privilege mode.

```
login as: root
root@mcctb's password:
Last login: Fri Jan  8 21:33:00 2017 from host.domain.com
```

2. Install the MariaDB server:

```
[root@mcctb ~]# yum install mariadb-server.x86_64
```

```
[root@mcctb ~]# yum install mariadb-server.x86_64
Loaded plugins: fastestmirror, langpacks
...
...

=====
===
Package                Arch   Version           Repository
Size
=====
Installing:
mariadb-server         x86_64  1:5.5.56-2.el7   base
11 M
```

```
Installing for dependencies:
```

```
Transaction Summary
```

```
=====
===
```

```
Install 1 Package (+8 Dependent packages)
Upgrade          ( 1 Dependent package)
```

```
Total download size: 22 M
```

```
Is this ok [y/d/N]: y
```

```
Downloading packages:
```

```
No Presto metadata available for base warning:
```

```
/var/cache/yum/x86_64/7/base/packages/mariadb-libs-5.5.56-2.e17.x86_64.rpm:
```

```
Header V3 RSA/SHA256 Signature,
```

```
key ID f4a80eb5: NOKEY] 1.4 MB/s | 3.3 MB 00:00:13 ETA
```

```
Public key for mariadb-libs-5.5.56-2.e17.x86_64.rpm is not installed
```

```
(1/10): mariadb-libs-5.5.56-2.e17.x86_64.rpm | 757 kB 00:00:01
```

```
..
```

```
..
```

```
(10/10): perl-Net-Daemon-0.48-5.e17.noarch.rpm | 51 kB 00:00:01
```

```
-----
-----
```

```
Installed:
```

```
  mariadb-server.x86_64 1:5.5.56-2.e17
```

```
Dependency Installed:
```

```
  mariadb.x86_64 1:5.5.56-2.e17
```

```
  perl-Compress-Raw-Bzip2.x86_64 0:2.061-3.e17
```

```
  perl-Compress-Raw-Zlib.x86_64 1:2.061-4.e17
```

```
  perl-DBD-MySQL.x86_64 0:4.023-5.e17
```

```
  perl-DBI.x86_64 0:1.627-4.e17
```

```
  perl-IO-Compress.noarch 0:2.061-2.e17
```

```
  perl-Net-Daemon.noarch 0:0.48-5.e17
```

```
  perl-PlRPC.noarch 0:0.2020-14.e17
```

```
Dependency Updated:
```

```
  mariadb-libs.x86_64 1:5.5.56-2.e17
```

```
Complete!
```

3. Start MariaDB server:

```
[root@mcctb ~]# systemctl start mariadb
```

4. Verify that the MariaDB server has started:

```
[root@mcctb ~]# systemctl status mariadb
```

```
[root@mcctb ~]# systemctl status mariadb
mariadb.service - MariaDB database server
...
Nov 08 21:28:59 mcctb systemd[1]: Starting MariaDB database server...
...
Nov 08 21:29:01 mcctb systemd[1]: Started MariaDB database server.
```

5. Configure the security and password settings:



When you are prompted for the root password, leave it empty and press enter to continue to configure the security and password settings.

```
[root@mcctb ~]# mysql_secure_installation
```

```
root@localhost systemd]# mysql_secure_installation
```

```
NOTE: RUNNING ALL PARTS OF THIS SCRIPT IS RECOMMENDED FOR ALL MariaDB
SERVERS IN PRODUCTION USE! PLEASE READ EACH STEP CAREFULLY!
```

In order to log into MariaDB to secure it, we'll need the current password for the root user. If you've just installed MariaDB, and you haven't set the root password yet, the password will be blank, so you should just press enter here.

```
Enter current password for root (enter for none):
OK, successfully used password, moving on...
```

Setting the root password ensures that nobody can log into the MariaDB root user without the proper authorisation.

```
Set root password? [Y/n] y
```

```
New password:
```

```
Re-enter new password:
```

```
Password updated successfully!
Reloading privilege tables..
... Success!
```

By default, a MariaDB installation has an anonymous user, allowing anyone to log into MariaDB without having to have a user account created for them. This is intended only for testing, and to make the installation go a bit smoother. You should remove them before moving into a

production environment.

Remove anonymous users? [Y/n] y

... Success!

Normally, root should only be allowed to connect from 'localhost'. This ensures that someone cannot guess at the root password from the network.

Disallow root login remotely? [Y/n] y

... Success!

By default, MariaDB comes with a database named 'test' that anyone can access. This is also intended only for testing, and should be removed before moving into a production environment.

Remove test database and access to it? [Y/n] y

- Dropping test database...

... Success!

- Removing privileges on test database...

... Success!

Reloading the privilege tables will ensure that all changes made so far will take effect immediately.

Reload privilege tables now? [Y/n]

... Success!

Cleaning up...

All done! If you've completed all of the above steps, your MariaDB installation should now be secure.

Thanks for using MariaDB!

Enable the autostart setting for the MariaDB server

You should verify that the autostart feature is turned on for the MariaDB server. If you do not enable the autostart feature, and the system on which the MetroCluster Tiebreaker software resides has to reboot, then the Tiebreaker software continues running, but the MariaDB service cannot be restarted and configuration changes cannot be made.

Steps

1. Enable the autostart service:

```
[root@mcctb ~]# systemctl enable mariadb.service
```

2. Verify that MariaDB is enabled to autostart when booted:

```
[root@mcctb ~]# systemctl list-unit-files mariadb.service
```

UNIT FILE	State
-----	-----
mariadb.service	enabled

Install or upgrade to Tiebreaker 1.5

Perform a new installation or upgrade to Tiebreaker 1.5 on your host Linux operating system to monitor MetroCluster configurations.

About this task

- Your storage system must be running a supported version of ONTAP. See the [Software requirements](#) table for more details.
- You must have installed OpenJDK by using the `yum install java-x.x.x-openjdk` command. Tiebreaker 1.5 and later supports OpenJDK 17, 18, or 19.
- You can install MetroCluster Tiebreaker as a non-root user with sufficient administrative privileges to perform the Tiebreaker installation, create tables and users, and set the user password.

Steps

1. Download the MetroCluster Tiebreaker software and the `MetroCluster_Tiebreaker_RPM_GPG` key.



The `MetroCluster_Tiebreaker_RPM_GPG` key is available to download from the same page that you download the software package for Tiebreaker 1.5 on the [NetApp Support Site](#).

[MetroCluster Tiebreaker \(Downloads\) - NetApp Support Site](#)

2. Log in to the host as the root user.
3. Create a non-root user and the `mcctbgrp` group.
 - a. Create a non-root user and set the password.

The following example commands create a non-root user named `mcctbuser1`:

```
[root@mcctb ~]# useradd mcctbuser1
[root@mcctb ~]# passwd mcctbuser1
Changing password for user mcctbuser1.
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

- b. Create a group named `mcctbgrp`:

```
[root@mcctb ~]# groupadd mcctbgrp
```

- c. Add the non-root user you created to the `mcctbgrp` group.

The following command adds `mcctbuser1` to the `mcctbgrp` group:

```
[root@mcctb ~]# usermod -a -G mcctbgrp mcctbuser1
```

4. Verify the RPM file.

Run the following substeps from the directory containing the RPM key.

- a. Download and import the RPM key file:

```
[root@mcctb ~]# rpm --import MetroCluster_Tiebreaker_RPM_GPG.key
```

- b. Verify that the correct key was imported by checking the fingerprint.

The following example shows a correct key fingerprint:

```
root@mcctb:~/signing/mcctb-rpms# gpg --show-keys --with-fingerprint
MetroCluster_Tiebreaker_RPM_GPG.key
pub   rsa3072 2022-11-17 [SCEA] [expires: 2025-11-16]
       65AC 1562 E28A 1497 7BBD 7251 2855 EB02 3E77 FAE5
uid           MCCTB-RPM (mcctb RPM production signing)
<mcctb-rpm@netapp.com>
```

- c. Verify the signature: `rpm --checksig NetApp-MetroCluster-Tiebreaker-Software-1.5-1.x86_64.rpm`

```
NetApp-MetroCluster-Tiebreaker-Software-1.5-1.x86_64.rpm: digests OK
```



You must only proceed with installation after you have successfully verified the signature.

5. Install or upgrade the Tiebreaker software:



You can only upgrade to Tiebreaker version 1.5 when you are upgrading from Tiebreaker version 1.4. Upgrading from earlier versions to Tiebreaker 1.5 is not supported.

Select the correct procedure depending on whether you're performing a new installation or upgrading an existing installation.

Perform a new installation

- a. Retrieve and record the absolute path for Java:

```
[root@mcctb ~]# readlink -f /usr/bin/java
/usr/lib/jvm/java-19-openjdk-19.0.0.0.36-
2.rolling.el8.x86_64/bin/java
```

- b. Run the following command:

```
rpm -ivh NetApp-MetroCluster-Tiebreaker-Software-1.5-1.x86_64.rpm
```

The system displays the following output for a successful installation:



When prompted during the installation, provide the non-root user that you previously created and assigned to the `mcctbgrp` group.

Example

```
Verifying...
##### [100%]
Preparing...
##### [100%]
Updating / installing...
  1:NetApp-MetroCluster-Tiebreaker-
So##### [100%]
Enter the absolute path for Java : /usr/lib/jvm/java-19-
openjdk-19.0.0.0.36-2.rolling.el8.x86_64/bin/java
Verifying if Java exists...
Found Java. Proceeding with the installation.
Enter host user account to use for the installation:
mcctbuser1
User account mcctbuser1 found. Proceeding with the
installation
Enter database user name:
root
Please enter database password for root
Enter password:
Sealed          false
Do you wish to auto unseal vault(y/n)?y
Enter the key1:
Enter the key2:
Enter the key3:
Success! Uploaded policy: mcctb-policy
Error enabling approle auth: Error making API request.
URL: POST http://127.0.0.1:8200/v1/sys/auth/approle
Code: 400. Errors:
* path is already in use at approle/
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/approle/role/mcctb-app
Password updated successfully in the vault.
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
Created symlink /etc/systemd/system/multi-
user.target.wants/netapp-metrocluster-tiebreaker-
software.service → /etc/systemd/system/netapp-metrocluster-
tiebreaker-software.service.
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
```

Successfully installed NetApp MetroCluster Tiebreaker software version 1.5.

Upgrading an existing installation

- a. Verify that a supported version of OpenJDK is installed and is the current Java version located on the host.



For upgrades to Tiebreaker 1.5, you must install either OpenJDK version 17, 18, or 19.

```
[root@mcctb ~]# readlink -f /usr/bin/java
/usr/lib/jvm/java-19-openjdk-19.0.0.0.36-
2.rolling.el8.x86_64/bin/java
```

- b. Verify the Vault service is unsealed and running: `vault status`

```
[root@mcctb ~]# vault status
Key          Value
---          -
Seal Type    shamir
Initialized   true
Sealed       false
Total Shares  5
Threshold    3
Version      1.12.2
Build Date   2022-11-23T12:53:46Z
Storage Type  file
Cluster Name  vault
Cluster ID    <cluster_id>
HA Enabled    false
```

- c. Upgrade the Tiebreaker software.

```
[root@mcctb ~]# rpm -Uvh NetApp-MetroCluster-Tiebreaker-Software-
1.5-1.x86_64.rpm
```

The system displays the following output for a successful upgrade:

Example

```
Verifying...
##### [100%]
Preparing...
##### [100%]
Updating / installing...
  1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]

Enter the absolute path for Java : /usr/lib/jvm/java-19-
openjdk-19.0.0.0.36-2.rolling.el8.x86_64/bin/java
Verifying if Java exists...
Found Java. Proceeding with the installation.
Enter host user account to use for the installation:
mcctbuser1
User account mcctbuser1 found. Proceeding with the
installation
Sealed          false
Do you wish to auto unseal vault(y/n)?y
Enter the key1:
Enter the key2:
Enter the key3:
Success! Uploaded policy: mcctb-policy
Error enabling approle auth: Error making API request.
URL: POST http://127.0.0.1:8200/v1/sys/auth/approle
Code: 400. Errors:
* path is already in use at approle/
Success! Enabled the kv secrets engine at: mcctb/
Success! Data written to: auth/approle/role/mcctb-app
Enter database user name : root
Please enter database password for root
Enter password:
Password updated successfully in the database.
Password updated successfully in the vault.
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable
netapp-metrocluster-tiebreaker-software
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Successfully upgraded NetApp MetroCluster Tiebreaker software
to version 1.5.
Cleaning up / removing...
```

```
2:NetApp-MetroCluster-Tiebreaker-  
So##### [100%]
```



If you enter the wrong MySQL root password, the Tiebreaker software indicates that it was installed successfully, but displays "Access denied" messages. To resolve the issue, you must uninstall the Tiebreaker software by using the `rpm -e` command, and then reinstall the software by using the correct MySQL root password.

6. Check the Tiebreaker connectivity to the MetroCluster software by opening an SSH connection from the Tiebreaker host to each of the node management LIFs and cluster management LIFs.

Related information

[NetApp Support](#)

Install Tiebreaker 1.4

Install MetroCluster Tiebreaker 1.4 dependencies

Depending on your host Linux operating system, install a MySQL or MariaDB server before installing or upgrading the Tiebreaker software.

Steps

1. [Install JDK](#).
2. Install MySQL or MariaDB server:

If the Linux host is	Then...
Red Hat Enterprise Linux 7/CentOS 7	Install MySQL Server 5.5.30 or later and 5.6.x versions on Red Hat Enterprise Linux 7 or CentOS 7
Red Hat Enterprise Linux 8	Install MariaDB server on Red Hat Enterprise Linux 8

Install JDK

You must install JDK on your host system before installing or upgrading the Tiebreaker software. Tiebreaker 1.4 and earlier supports JDK 1.8.0. (JRE 8).

Steps

1. Log in as a "root" user.

```
login as: root  
root@mcctb's password:  
Last login: Fri Jan 8 21:33:00 2017 from host.domain.com
```

2. Install JDK 1.8.0:

```
yum install java-1.8.0-openjdk.x86_64
```

```
[root@mcctb ~]# yum install java-1.8.0-openjdk.x86_64
Loaded plugins: fastestmirror, langpacks
Loading mirror speeds from cached hostfile
... shortened....
Dependencies Resolved

=====
Package                Arch    Version                               Repository    Size
=====
Installing:
  java-1.8.0-openjdk    x86_64  1:1.8.0.144-0.b01.e17_4             updates      238 k
  ..
  ..
Transaction Summary
=====
Install 1 Package (+ 4 Dependent packages)

Total download size: 34 M
Is this ok [y/d/N]: y

Installed:
java-1.8.0-openjdk.x86_64 1:1.8.0.144-0.b01.e17_4
Complete!
```

Install MySQL Server 5.5.30 or later and 5.6.x versions on Red Hat Enterprise Linux 7 or CentOS 7

You must install MySQL Server 5.5.30 or later and 5.6.x version on your host system before installing or upgrading the Tiebreaker software. For Red Hat Enterprise Linux 8, [Install the MariaDB server](#).

Steps

1. Log in as a root user.

```
login as: root
root@mcctb's password:
Last login: Fri Jan  8 21:33:00 2016 from host.domain.com
```

2. Add the MySQL repository to your host system:

```
[root@mcctb ~]# yum localinstall https://dev.mysql.com/get/mysql57-community-release-el6-11.noarch.rpm
```

```

Loaded plugins: product-id, refresh-packagekit, security, subscription-
manager
Setting up Local Package Process
Examining /var/tmp/yum-root-LLUw0r/mysql-community-release-el6-
5.noarch.rpm: mysql-community-release-el6-5.noarch
Marking /var/tmp/yum-root-LLUw0r/mysql-community-release-el6-
5.noarch.rpm to be installed
Resolving Dependencies
--> Running transaction check
---> Package mysql-community-release.noarch 0:el6-5 will be installed
--> Finished Dependency Resolution
Dependencies Resolved

=====
=====
Package                Arch    Version
                        Repository

Size
=====
=====
Installing:
mysql-community-release
                        noarch el6-5 /mysql-community-release-el6-
5.noarch 4.3 k
Transaction Summary
=====
=====
Install                1 Package(s)
Total size: 4.3 k
Installed size: 4.3 k
Is this ok [y/N]: y
Downloading Packages:
Running rpm_check_debug
Running Transaction Test
Transaction Test Succeeded
Running Transaction
  Installing : mysql-community-release-el6-5.noarch
1/1
  Verifying  : mysql-community-release-el6-5.noarch
1/1
Installed:
  mysql-community-release.noarch 0:el6-5
Complete!

```

3. Disable the MySQL 57 repository:

```
[root@mcctb ~]# yum-config-manager --disable mysql57-community
```

4. Enable the MySQL 56 repository:

```
[root@mcctb ~]# yum-config-manager --enable mysql56-community
```

5. Enable the repository:

```
[root@mcctb ~]# yum repolist enabled | grep "mysql.-community."
```

```
mysql-connectors-community           MySQL Connectors Community
21
mysql-tools-community                MySQL Tools Community
35
mysql56-community                    MySQL 5.6 Community Server
231
```

6. Install the MySQL Community server:

```
[root@mcctb ~]# yum install mysql-community-server
```

```
Loaded plugins: product-id, refresh-packagekit, security, subscription-
manager
This system is not registered to Red Hat Subscription Management. You
can use subscription-manager
to register.
Setting up Install Process
Resolving Dependencies
--> Running transaction check
.....Output truncated.....
---> Package mysql-community-libs-compat.x86_64 0:5.6.29-2.el6 will be
obsoleting
--> Finished Dependency Resolution
Dependencies Resolved

=====
=====
Package                               Arch    Version           Repository
Size
=====
Installing:
mysql-community-client                 x86_64  5.6.29-2.el6     mysql56-community
18 M
    replacing mysql.x86_64 5.1.71-1.el6
mysql-community-libs                   x86_64  5.6.29-2.el6     mysql56-community
1.9 M
```

```
replacing mysql-libs.x86_64 5.1.71-1.el6
mysql-community-libs-compat x86_64 5.6.29-2.el6 mysql56-community
1.6 M
replacing mysql-libs.x86_64 5.1.71-1.el6
mysql-community-server x86_64 5.6.29-2.el6 mysql56-community
53 M
replacing mysql-server.x86_64 5.1.71-1.el6
Installing for dependencies:
mysql-community-common x86_64 5.6.29-2.el6 mysql56-community
308 k
```

Transaction Summary

=====

=====

Install 5 Package(s)

Total download size: 74 M

Is this ok [y/N]: y

Downloading Packages:

```
(1/5): mysql-community-client-5.6.29-2.el6.x86_64.rpm | 18 MB
00:28
(2/5): mysql-community-common-5.6.29-2.el6.x86_64.rpm | 308 kB
00:01
(3/5): mysql-community-libs-5.6.29-2.el6.x86_64.rpm | 1.9 MB
00:05
(4/5): mysql-community-libs-compat-5.6.29-2.el6.x86_64.rpm | 1.6 MB
00:05
(5/5): mysql-community-server-5.6.29-2.el6.x86_64.rpm | 53 MB
03:42
```

```
Total 289 kB/s | 74 MB
04:24
```

warning: rpmts_HdrFromFdno: Header V3 DSA/SHA1 Signature, key ID
<key_id> NOKEY

Retrieving key from file:/etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

Importing GPG key 0x5072E1F5:

 Userid : MySQL Release Engineering <mysql-build@oss.oracle.com>

 Package: mysql-community-release-el6-5.noarch

 (@/mysql-community-release-el6-5.noarch)

 From : file:/etc/pki/rpm-gpg/RPM-GPG-KEY-mysql

Is this ok [y/N]: y

Running rpm_check_debug

Running Transaction Test

Transaction Test Succeeded

Running Transaction

 Installing : mysql-community-common-5.6.29-2.el6.x86_64

```
....Output truncated....
```

```
1.el6.x86_64
```

```
7/8
```

```
Verifying : mysql-5.1.71-1.el6.x86_64
```

```
8/8
```

```
Installed:
```

```
mysql-community-client.x86_64 0:5.6.29-2.el6
```

```
mysql-community-libs.x86_64 0:5.6.29-2.el6
```

```
mysql-community-libs-compat.x86_64 0:5.6.29-2.el6
```

```
mysql-community-server.x86_64 0:5.6.29-2.el6
```

```
Dependency Installed:
```

```
mysql-community-common.x86_64 0:5.6.29-2.el6
```

```
Replaced:
```

```
mysql.x86_64 0:5.1.71-1.el6 mysql-libs.x86_64 0:5.1.71-1.el6
```

```
mysql-server.x86_64 0:5.1.71-1.el6
```

```
Complete!
```

7. Start MySQL server:

```
[root@mcctb ~]# service mysqld start
```

```
Initializing MySQL database: 2016-04-05 19:44:38 0 [Warning] TIMESTAMP
with implicit DEFAULT value is deprecated. Please use
--explicit_defaults_for_timestamp server option (see documentation
for more details).
2016-04-05 19:44:38 0 [Note] /usr/sbin/mysqld (mysqld 5.6.29)
starting as process 2487 ...
2016-04-05 19:44:38 2487 [Note] InnoDB: Using atomics to ref count
buffer pool pages
2016-04-05 19:44:38 2487 [Note] InnoDB: The InnoDB memory heap is
disabled
....Output truncated....
2016-04-05 19:44:42 2509 [Note] InnoDB: Shutdown completed; log sequence
number 1625987
```

PLEASE REMEMBER TO SET A PASSWORD FOR THE MySQL root USER!
To do so, start the server, then issue the following commands:

```
/usr/bin/mysqladmin -u root password 'new-password'
/usr/bin/mysqladmin -u root -h mcctb password 'new-password'
```

Alternatively, you can run:

```
/usr/bin/mysql_secure_installation
```

which will also give you the option of removing the test
databases and anonymous user created by default. This is
strongly recommended for production servers.

.....Output truncated.....

```
WARNING: Default config file /etc/my.cnf exists on the system
This file will be read by default by the MySQL server
If you do not want to use this, either remove it, or use the
--defaults-file argument to mysqld_safe when starting the server
```

```
Starting mysqld: [ OK ]
```

8. Confirm that MySQL server is running:

```
[root@mcctb ~]# service mysqld status
```

```
mysqld (pid 2739) is running...
```

9. Configure security and password settings:

```
[root@mcctb ~]# mysql_secure_installation
```

NOTE: RUNNING ALL PARTS OF THIS SCRIPT IS RECOMMENDED FOR ALL MySQL
SERVERS IN PRODUCTION USE! PLEASE READ EACH STEP CAREFULLY!

In order to log into MySQL to secure it, we'll need the current password for the root user. If you've just installed MySQL, and you haven't set the root password yet, the password will be blank, so you should just press enter here.

Enter current password for root (enter for none): <== on default
install

hit enter here

OK, successfully used password, moving on...

Setting the root password ensures that nobody can log into the MySQL root user without the proper authorization.

Set root password? [Y/n] y

New password:

Re-enter new password:

Password updated successfully!

Reloading privilege tables..

... Success!

By default, a MySQL installation has an anonymous user, allowing anyone to log into MySQL without having to have a user account created for them. This is intended only for testing, and to make the installation go a bit smoother. You should remove them before moving into a production environment.

Remove anonymous users? [Y/n] y

... Success!

Normally, root should only be allowed to connect from 'localhost'. This ensures that someone cannot guess at the root password from the network.

Disallow root login remotely? [Y/n] y

... Success!

By default, MySQL comes with a database named 'test' that anyone can access. This is also intended only for testing, and should be removed before moving into a production environment.

Remove test database and access to it? [Y/n] y

- Dropping test database...

ERROR 1008 (HY000) at line 1: Can't drop database 'test';

```
database doesn't exist
... Failed! Not critical, keep moving...
- Removing privileges on test database...
... Success!

Reloading the privilege tables will ensure that all changes made so far
will take effect immediately.

Reload privilege tables now? [Y/n] y
... Success!

All done! If you've completed all of the above steps, your MySQL
installation should now be secure.

Thanks for using MySQL!

Cleaning up...
```

10. Verify that the MySQL login is working:

```
[root@mcctb ~]# mysql -u root -p
```

```
Enter password: <configured_password>
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 17
Server version: 5.6.29 MySQL Community Server (GPL)

Copyright (c) 2000, 2016, Oracle and/or its affiliates. All rights
reserved.

Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input
statement.
mysql>
```

When the MySQL login is working as expected, the output ends at the `mysql>` prompt.

Enable the MySQL autostart setting

You should verify that the autostart feature is turned on for the MySQL daemon. Turning on the MySQL daemon automatically restarts MySQL if the system on which the MetroCluster Tiebreaker software resides reboots. If the MySQL daemon is not running, the Tiebreaker software continues running, but it cannot be restarted and configuration changes cannot be made.

Step

1. Verify that MySQL is enabled to autostart when booted:

```
[root@mcctb ~]# systemctl list-unit-files mysqld.service
```

```
UNIT FILE           State
-----
mysqld.service     enabled
```

If MySQL is not enabled to autostart when booted, see the MySQL documentation to enable the autostart feature for your installation.

Install MariaDB server on Red Hat Enterprise Linux 8

You must install MariaDB server on your host system before installing or upgrading the Tiebreaker software. For Red Hat Enterprise Linux 7 or CentOS 7, [Install MySQL Server](#).

Before you begin

Your host system must be running on Red Hat Enterprise Linux (RHEL) 8.

Steps

1. Log in as a root user.

```
login as: root
root@mcctb's password:
Last login: Fri Jan  8 21:33:00 2017 from host.domain.com
```

2. Install the MariaDB server:

```
[root@mcctb ~]# yum install mariadb-server.x86_64
```

```
[root@mcctb ~]# yum install mariadb-server.x86_64
Loaded plugins: fastestmirror, langpacks
...
...

=====
===
Package                Arch   Version           Repository
Size
=====
Installing:
mariadb-server         x86_64  1:5.5.56-2.el7    base
11 M
```

```
Installing for dependencies:
```

```
Transaction Summary
```

```
=====
```

```
Install 1 Package (+8 Dependent packages)
Upgrade ( 1 Dependent package)
```

```
Total download size: 22 M
```

```
Is this ok [y/d/N]: y
```

```
Downloading packages:
```

```
No Presto metadata available for base warning:
```

```
/var/cache/yum/x86_64/7/base/packages/mariadb-libs-5.5.56-2.e17.x86_64.rpm:
```

```
Header V3 RSA/SHA256 Signature,
```

```
key ID f4a80eb5: NOKEY] 1.4 MB/s | 3.3 MB 00:00:13 ETA
```

```
Public key for mariadb-libs-5.5.56-2.e17.x86_64.rpm is not installed
```

```
(1/10): mariadb-libs-5.5.56-2.e17.x86_64.rpm | 757 kB 00:00:01
```

```
..
```

```
..
```

```
(10/10): perl-Net-Daemon-0.48-5.e17.noarch.rpm | 51 kB 00:00:01
```

```
-----
```

```
Installed:
```

```
  mariadb-server.x86_64 1:5.5.56-2.e17
```

```
Dependency Installed:
```

```
  mariadb.x86_64 1:5.5.56-2.e17
```

```
  perl-Compress-Raw-Bzip2.x86_64 0:2.061-3.e17
```

```
  perl-Compress-Raw-Zlib.x86_64 1:2.061-4.e17
```

```
  perl-DBD-MySQL.x86_64 0:4.023-5.e17
```

```
  perl-DBI.x86_64 0:1.627-4.e17
```

```
  perl-IO-Compress.noarch 0:2.061-2.e17
```

```
  perl-Net-Daemon.noarch 0:0.48-5.e17
```

```
  perl-PlRPC.noarch 0:0.2020-14.e17
```

```
Dependency Updated:
```

```
  mariadb-libs.x86_64 1:5.5.56-2.e17
```

```
Complete!
```

3. Start MariaDB server:

```
[root@mcctb ~]# systemctl start mariadb
```

4. Verify that the MariaDB server has started:

```
[root@mcctb ~]# systemctl status mariadb
```

```
[root@mcctb ~]# systemctl status mariadb
mariadb.service - MariaDB database server
...
Nov 08 21:28:59 mcctb systemd[1]: Starting MariaDB database server...
...
Nov 08 21:29:01 mcctb systemd[1]: Started MariaDB database server.
```

5. Configure the security and password settings:



When you are prompted for the root password, leave it empty and press enter to continue to configure the security and password settings.

```
[root@mcctb ~]# mysql_secure_installation
```

```
root@localhost systemd]# mysql_secure_installation
```

```
NOTE: RUNNING ALL PARTS OF THIS SCRIPT IS RECOMMENDED FOR ALL MariaDB
SERVERS IN PRODUCTION USE! PLEASE READ EACH STEP CAREFULLY!
```

In order to log into MariaDB to secure it, we'll need the current password for the root user. If you've just installed MariaDB, and you haven't set the root password yet, the password will be blank, so you should just press enter here.

```
Enter current password for root (enter for none):
OK, successfully used password, moving on...
```

Setting the root password ensures that nobody can log into the MariaDB root user without the proper authorisation.

```
Set root password? [Y/n] y
```

```
New password:
```

```
Re-enter new password:
```

```
Password updated successfully!
Reloading privilege tables..
... Success!
```

By default, a MariaDB installation has an anonymous user, allowing anyone to log into MariaDB without having to have a user account created for them. This is intended only for testing, and to make the installation go a bit smoother. You should remove them before moving into a

production environment.

Remove anonymous users? [Y/n] y

... Success!

Normally, root should only be allowed to connect from 'localhost'. This ensures that someone cannot guess at the root password from the network.

Disallow root login remotely? [Y/n] y

... Success!

By default, MariaDB comes with a database named 'test' that anyone can access. This is also intended only for testing, and should be removed before moving into a production environment.

Remove test database and access to it? [Y/n] y

- Dropping test database...

... Success!

- Removing privileges on test database...

... Success!

Reloading the privilege tables will ensure that all changes made so far will take effect immediately.

Reload privilege tables now? [Y/n]

... Success!

Cleaning up...

All done! If you've completed all of the above steps, your MariaDB installation should now be secure.

Thanks for using MariaDB!

Enable the autostart setting for the MariaDB server

You should verify that the autostart feature is turned on for the MariaDB server. If you do not enable the autostart feature, and the system on which the MetroCluster Tiebreaker software resides has to reboot, then the Tiebreaker software continues running, but the MariaDB service cannot be restarted and configuration changes cannot be made.

Steps

1. Enable the autostart service:

```
[root@mcctb ~]# systemctl enable mariadb.service
```

2. Verify that MariaDB is enabled to autostart when booted:

```
[root@mcctb ~]# systemctl list-unit-files mariadb.service
```

UNIT FILE	State
-----	-----
mariadb.service	enabled

Install or upgrade to Tiebreaker 1.4

Perform a new installation or upgrade to Tiebreaker 1.4 on your host Linux operating system to monitor MetroCluster configurations.

About this task

- Your storage system must be running a supported version of ONTAP. See the [Software requirements](#) table for more details.
- You must have installed OpenJDK by using the `yum install java-x.x.x-openjdk` command. Tiebreaker 1.4 and earlier supports JDK 1.8.0 (JRE 8).

Steps

1. Download the MetroCluster Tiebreaker software.

[MetroCluster Tiebreaker \(Downloads\) - NetApp Support Site](#)

2. Log in to the host as the root user.

3. Install or upgrade the Tiebreaker software:

Select the correct procedure depending on whether you're performing a new installation or upgrading an existing installation.

Perform a new installation

- a. Install the Tiebreaker software by running the :

```
rpm -ivh NetApp-MetroCluster-Tiebreaker-Software-1.4-1.x86_64.rpm
```

The system displays the following output for a successful installation:

```
Verifying...
##### [100%]
Preparing...
##### [100%]
Updating / installing...
   1:NetApp-MetroCluster-Tiebreaker-
So##### [100%]
Post installation start Fri Apr  5 02:28:09 EDT 2024
Enter MetroCluster Tiebreaker user password:

Please enter mysql root password when prompted
Enter password:
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable netapp-
metrocluster-tiebreaker-software
Created symlink /etc/systemd/system/multi-
user.target.wants/netapp-metrocluster-tiebreaker-software.service
→ /etc/systemd/system/netapp-metrocluster-tiebreaker-
software.service.
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Enabled autostart of NetApp MetroCluster Tiebreaker software
daemon during boot
Created symbolic link for NetApp MetroCluster Tiebreaker software
CLI
Post installation end Fri Apr  5 02:28:22 EDT 2024
Successfully installed NetApp MetroCluster Tiebreaker software
version 1.4.
```

Upgrade an existing installation

- a. Upgrade the Tiebreaker software.

```
[root@mcctb ~]# rpm -Uvh NetApp-MetroCluster-Tiebreaker-Software-1.4-1.x86_64.rpm
```

The system displays the following output for a successful upgrade:

```
Verifying...
##### [100%]
Preparing...
##### [100%]
Upgrading NetApp MetroCluster Tiebreaker software....
Stopping NetApp MetroCluster Tiebreaker software services before
upgrade.
Updating / installing...
 1:NetApp-MetroCluster-Tiebreaker-
So##### [ 50%]
Post installation start Mon Apr  8 06:29:51 EDT 2024
Synchronizing state of netapp-metrocluster-tiebreaker-
software.service with SysV service script with
/usr/lib/systemd/systemd-sysv-install.
Executing: /usr/lib/systemd/systemd-sysv-install enable netapp-
metrocluster-tiebreaker-software
Attempting to start NetApp MetroCluster Tiebreaker software
services
Started NetApp MetroCluster Tiebreaker software services
Enabled autostart of NetApp MetroCluster Tiebreaker software
daemon during boot
Created symbolic link for NetApp MetroCluster Tiebreaker software
CLI
Post upgrade end Mon Apr  8 06:29:51 EDT 2024
Successfully upgraded NetApp MetroCluster Tiebreaker software to
version 1.4.
Cleaning up / removing...
 2:NetApp-MetroCluster-Tiebreaker-
So##### [100%]
```



If you enter the wrong MySQL root password, the Tiebreaker software indicates that it was installed successfully, but displays "Access denied" messages. To resolve the issue, you must uninstall the Tiebreaker software by using the `rpm -e` command, and then reinstall the software by using the correct MySQL root password.

4. Check the Tiebreaker connectivity to the MetroCluster software by opening an SSH connection from the Tiebreaker host to each of the node management LIFs and cluster management LIFs.

Related information

Upgrade the host where the Tiebreaker monitor is running

You might need to upgrade the host that the Tiebreaker monitor is running on.

Steps

1. Uninstall the Tiebreaker software:

```
rpm -e NetApp-MetroCluster-Tiebreaker-Software
```

2. Upgrade the host. Refer to your host OS documentation for details.
3. Reinstall the Tiebreaker software.

Perform a fresh installation of Tiebreaker by following the steps in [Install the Tiebreaker software](#).

Configure the MetroCluster Tiebreaker software

After installation of the Tiebreaker software, you can add or modify MetroCluster configurations, or remove them from the Tiebreaker software.

Launch the Tiebreaker software CLI

After installing the Tiebreaker software, you must launch its CLI to configure the software.

1. Launch the CLI from the prompt of the host on which you installed the software:

```
netapp-metrocluster-tiebreaker-software-cli
```

2. After installation and during the first startup, enter the password for the Tiebreaker user to access the database.
This is the password that you specified for the database user during installation.

Add MetroCluster configurations

After installing the NetApp MetroCluster Tiebreaker software, you can add more MetroCluster configurations, one at a time.

About this task

- You must have installed the MetroCluster configuration in an ONTAP environment and enabled the settings in the software.
- The steps and expected command output depends on the version of Tiebreaker that you are running.

Tiebreaker 1.5 or earlier

Steps

1. Use the Tiebreaker command-line interface (CLI) monitor add command to add MetroCluster configurations.

If you are using the host name, it must be the fully qualified domain name (FQDN).

The following example shows the configuration of cluster_A:

Example

```
NetApp MetroCluster Tiebreaker :> monitor add wizard
Enter monitor Name: <monitor_name>
Enter Cluster IP Address: <cluster_ip_value>
Enter Cluster Username: admin
Enter Cluster Password:
Enter Cluster IP Address: <peer_cluster_ip_value>
Enter Peer Cluster Username: admin
Enter Peer Cluster Password:
Successfully added monitor to NetApp MetroCluster Tiebreaker
software.
```

2. Confirm that the MetroCluster configuration was added properly by using the Tiebreaker CLI monitor show -status command.

```
NetApp MetroCluster Tiebreaker :> monitor show -status
```

3. Disable the observer mode for the Tiebreaker software to automatically initiate a switchover after it detects a site failure:

```
monitor modify -monitor-name <monitor_name> -observer-mode false
```

```
NetApp MetroCluster Tiebreaker :> monitor modify -monitor-name 8pack
-observer-mode false
Warning: If you are turning observer-mode to false, make sure to
review the 'risks and limitations'
as described in the MetroCluster Tiebreaker installation and
configuration.
Are you sure you want to enable automatic switchover capability for
monitor "8pack"? [Y/N]: y
```

Tiebreaker 1.6 or later

Steps

1. Use the Tiebreaker command-line interface (CLI) monitor add command to add MetroCluster configurations.

If you are using the host name, it must be the fully qualified domain name (FQDN).

The following example shows the configuration of cluster_A:

Example

```
NetApp MetroCluster Tiebreaker :> monitor add wizard
Enter Monitor Name: cluster_A
Enter Cluster IP Address: <cluster_ip_value>
Enter Cluster Username: admin
Enter Cluster Password:
Enter Peer Cluster IP Address: <peer_cluster_ip_value>
Enter Peer Cluster Username: admin
Enter Peer Cluster Password:
```

NOTE: Before enabling automatic switchover capability, make sure to review the 'risks and limitations' as described in the MetroCluster Tiebreaker Installation and Configuration Guide.

```
Do you want to enable automatic switchover capability for
monitor(Y/N): y
Successfully added monitor to NetApp MetroCluster Tiebreaker
software.
Verifying SSL certificate chain from cluster_A...
```

```
=====
Warning missing SSL certificates
=====
```

```
Cluster: cluster_A

IP Address: <cluster_ip_value>
```

Result:

The MetroCluster Tiebreaker is unable to verify the SSL certificate chain.

Recommended Actions:

Run the following command to identify missing certificates:

```
monitor switchover-simulate
```

Import any missing certificates as indicated by the command output.

For detailed instructions, please refer to the MetroCluster Tiebreaker documentation, or contact NetApp Support for assistance.

Note:

Missing certificates will prevent the MetroCluster Tiebreaker from issuing a switchover request in the event of a site failure.

```
=====
```

```
Verifying SSL certificate chain from cluster_B...
```

```
SSL certificate chain is valid
```

2. Confirm that the MetroCluster configuration was added properly by using the Tiebreaker CLI `monitor show -status` command.

```
NetApp MetroCluster Tiebreaker :> monitor show -status
```

3. Import the missing certificates by performing the steps for your Tiebreaker version in [Import certificates](#).

Related information

[Risks and limitations of using MetroCluster Tiebreaker in active mode](#)

Import certificates

To enable seamless monitoring in Tiebreaker 1.6 or later, you need to import the server Secure Sockets Layer (SSL) certificate, intermediate certificate (if one exists), and the root certificate from ONTAP to the key store in the Java Virtual Machine (JVM).

About this task

- This task is required in Tiebreaker 1.6 or later.
- You perform this task after you have successfully added a MetroCluster configuration to Tiebreaker or if your certificates expire.
- In Tiebreaker 1.7 or later, you can perform a switchover simulation to check if you need to import a certificate. If the switchover simulation fails, you need to import certificates from ONTAP to the key store in the Java Virtual Machine (JVM).

Tiebreaker 1.7 or later

Steps

1. Run a switchover simulation to check if you need to import certificates.
 - a. Check the Tiebreaker monitoring status:

```
monitor show -status
```

Example

```
NetApp MetroCluster Tiebreaker :> monitor show -status
MetroCluster: A700
  Disaster: false
  Monitor State: Normal
  Observer Mode: false
  Silent Period: 5
  Override Vetoes: false
  Cluster: ClusterA_siteA (UUID:713e5ab2-b4e8-11f0-91aa-00a098ef36a2)
    Reachable: true
    Intersite Connectivity Available: true
      Node: node_A1 (UUID:9f6cecbf-b4e4-11f0-9d0f-00a098ef36a2)
        Reachable: true
        Intersite Connectivity Available: true
        State: normal
      Node: node_A2 (UUID:2719bb56-b4e7-11f0-996c-00a09897caa3)
        Reachable: true
        Intersite Connectivity Available: true
        State: normal
    Cluster: ClusterB_siteB (UUID:72839591-b4e8-11f0-b688-00a09897cb73)
      Reachable: true
      Intersite Connectivity Available: true
        Node: node_B1 (UUID:abfeab89-b4e4-11f0-a077-00a09897cb73)
          Reachable: true
          Intersite Connectivity Available: true
          State: normal
        Node: node_B2 (UUID:31e395bf-b4e7-11f0-bf99-00a09897cb2f)
          Reachable: true
          Intersite Connectivity Available: true
          State: normal
```

b. Trigger a switchover simulation:

```
monitor switchover-simulate -monitor-name <monitor_name> -cluster  
<cluster_name>
```

The command returns the following output if you need to import certificates to the JVM:

```
Failed to trigger Switchover Simulation. Please check  
Metrocluster Tiebreaker logs for further information or contact  
NetApp support.
```

2. Run the following command for each certificate you need to import (SSL server, intermediate, or root).

```
/opt/netapp/java/bin/keytool -import -trustcacerts -file  
<certificate_file_name> -keystore "/opt/netapp/java/lib/security/cacerts"  
-alias <certificate>
```

- The <certificate_file_name> value specifies the file name of the certificate that you want to import.
- The -alias <certificate> value specifies the name that you want to store the certificate under after it is imported to the JVM.

The following example shows how to import a root certificate with the file name `root.crt` and an SSL server certificate with the file name `ssl_cert.crt`:

```
/opt/netapp/java/bin/keytool -import -trustcacerts -file root.crt  
-keystore "/opt/netapp/java/lib/security/cacerts" -alias root  
  
/opt/netapp/java/bin/keytool -import -trustcacerts -file  
ssl_cert.crt -keystore "/opt/netapp/java/lib/security/cacerts"  
-alias ssl_cert
```

3. Restart the Tiebreaker software:

```
systemctl restart netapp-metrocluster-tiebreaker-software
```

4. Perform the switchover simulation checks again:

a. Check the Tiebreaker monitoring status:

```
monitor show -status
```

Example

```
NetApp MetroCluster Tiebreaker :> monitor show -status
MetroCluster: A700
  Disaster: false
  Monitor State: Normal
  Observer Mode: false
  Silent Period: 5
  Override Vetoes: false
  Cluster: ClusterA_siteA(UUID:713e5ab2-b4e8-11f0-91aa-00a098ef36a2)
    Reachable: true
    Intersite Connectivity Available: true
      Node: node_A1(UUID:9f6cecbf-b4e4-11f0-9d0f-00a098ef36a2)
        Reachable: true
        Intersite Connectivity Available: true
        State: normal
      Node: node_A2(UUID:2719bb56-b4e7-11f0-996c-00a09897caa3)
        Reachable: true
        Intersite Connectivity Available: true
        State: normal
    Cluster: ClusterB_siteB(UUID:72839591-b4e8-11f0-b688-00a09897cb73)
      Reachable: true
      Intersite Connectivity Available: true
        Node: node_B1(UUID:abfeab89-b4e4-11f0-a077-00a09897cb73)
          Reachable: true
          Intersite Connectivity Available: true
          State: normal
        Node: node_B2(UUID:31e395bf-b4e7-11f0-bf99-00a09897cb2f)
          Reachable: true
          Intersite Connectivity Available: true
          State: normal
```

b. Trigger a switchover simulation:

```
monitor switchover-simulate -monitor-name <monitor_name> -cluster
<cluster_name>
```

```
Successfully triggered Switchover Simulation. Please check the
status of the Switchover Simulation on the ONTAP cluster using
command "metrocluster operation history show"
```

Tiebreaker 1.6 or 1.6P1

Steps

1. Import all certificates from ONTAP. Run the following command for each certificate you need to import (SSL server, intermediate, or root).

```
/opt/netapp/java/bin/keytool -import -trustcacerts -file
<certificate_file_name> -keystore "/opt/netapp/java/lib/security/cacerts"
-alias <certificate>
```

- The <certificate_file_name> value specifies the file name of the certificate that you want to import.
- The -alias <certificate> value specifies the name that you want to store the certificate under after it is imported to the JVM.

The following example shows how to import a root certificate with the file name `root.crt` and an SSL server certificate with the file name `ssl_cert.crt`:

```
/opt/netapp/java/bin/keytool -import -trustcacerts -file root.crt
-keystore "/opt/netapp/java/lib/security/cacerts" -alias root

/opt/netapp/java/bin/keytool -import -trustcacerts -file
ssl_cert.crt -keystore "/opt/netapp/java/lib/security/cacerts"
-alias ssl_cert
```

2. Restart the Tiebreaker software:

```
systemctl restart netapp-metrocluster-tiebreaker-software
```

Commands for modifying MetroCluster Tiebreaker configurations

You can modify the MetroCluster configuration whenever you need to change the settings.

The Tiebreaker CLI monitor modify command can be used with any of the following options. You can confirm your changes with the monitor show -status command.

Option	Description
-monitor-name	Name of the MetroCluster configuration
-enable-monitor	Enables and disables monitoring of the MetroCluster configuration

-silent-period	Period in seconds for which the MetroCluster Tiebreaker software waits to confirm a site failure after detection
-observer-mode	<p>Observer mode (true) provides monitoring only, and does not trigger a switchover if a site disaster occurs. Online mode (false) triggers a switchover if a site disaster occurs.</p> <ul style="list-style-type: none"> • How the Tiebreaker software detects site failure • Risks and limitations of using MetroCluster Tiebreaker in active mode

The following example changes the silent period for the configuration.

```
NetApp MetroCluster Tiebreaker :> monitor modify -monitor-name cluster_A
-silent-period 15
Successfully modified monitor in NetApp MetroCluster Tiebreaker
software.
```

The Tiebreaker CLI `debug` command can be used to change the logging mode.

Command	Description
debug status	Displays the status of the debug mode
debug enable	Enables the debug mode for logging
debug disable	Disables the debug mode for logging

In systems running Tiebreaker 1.4 and earlier, the Tiebreaker CLI `update-mcctb-password` command can be used to update the user password. This command is deprecated in Tiebreaker 1.5 and later.

Command	Description
update-mcctb-password	The user password is successfully updated

Remove MetroCluster configurations

You can remove the MetroCluster configuration that is being monitored by the Tiebreaker software when you no longer want to monitor a MetroCluster configuration.

1. Use the Tiebreaker CLI `monitor remove` command to remove the MetroCluster configuration.

In the following example, “cluster_A” is removed from the software:

```
NetApp MetroCluster Tiebreaker :> monitor remove -monitor-name cluster_A
Successfully removed monitor from NetApp MetroCluster Tiebreaker
software.
```

2. Confirm that the MetroCluster configuration is removed properly by using the Tiebreaker CLI `monitor show -status` command.

```
NetApp MetroCluster Tiebreaker :> monitor show -status
```

Configuring SNMP settings for Tiebreaker software

To use SNMP with the Tiebreaker software, you must configure the SNMP settings.

About this task

- Tiebreaker 1.6 only supports SNMPv3.
- Although Tiebreaker 1.5 and 1.4 support SNMPv1 and SNMPv3, NetApp strongly recommends that you configure SNMPv3 for optimum security.

Steps

1. Use the Tiebreaker CLI `snmp config wizard` command to add MetroCluster configurations.



Only one SNMP trap host is currently supported.

The `snmp config wizard` command response depends on the version of Tiebreaker you are running.

Tiebreaker 1.6

The following example shows the configuration of an SNMP receiver that supports SNMPv3 with an IP address of 192.0.2.255 and port number 162 for trap messages. The Tiebreaker software is ready to send traps to the SNMP receiver that you specified.



Tiebreaker 1.6 only supports SNMPv3

```
NetApp MetroCluster Tiebreaker :> snmp config wizard
Enter SNMP Host: 192.0.2.255
Enter SNMP Port: 162
Enter SNMP V3 Security Name: v3sec
Enter SNMP V3 Authentication password:
```

Tiebreaker 1.5 and 1.4

The following example shows the configuration of an SNMP receiver that supports SNMPv3 with an IP address of 192.0.2.255 and port number 162 for trap messages. The Tiebreaker software is ready to send traps to the SNMP receiver that you specified.

```
NetApp MetroCluster Tiebreaker :> snmp config wizard
Enter SNMP Version[V1/V3]: v3
Enter SNMP Host: 192.0.2.255
Enter SNMP Port: 162
Enter SNMP V3 Security Name: v3sec
Enter SNMP V3 Authentication password:
Enter SNMP V3 Privacy password:
Engine ID : 8000031504932eff571825192a6f1193b265e24593
Successfully added SNMP properties to NetApp MetroCluster Tiebreaker
software.
```



You should configure SNMPv3 because SNMPv1 is not secure. Verify that the default community string is **NOT** set to public.

2. Verify that the SNMP settings are configured:

```
snmp config test
```

The following example shows that the Tiebreaker software can send an SNMP trap for the event TEST_SNMP_CONFIG:

```
NetApp MetroCluster Tiebreaker :> snmp config test
Sending SNMP trap to localhost. Version : V3.
Successfully sent SNMP trap for event TEST_SNMP_CONFIG
NetApp MetroCluster Tiebreaker :>
```

Monitoring the MetroCluster configuration

MetroCluster Tiebreaker software automates the recovery process by enabling you to monitor the MetroCluster configuration status, evaluate SNMP events and traps that are sent to NetApp customer support, and view the status of monitoring operations.

Configuring AutoSupport

By default, AutoSupport messages are sent to NetApp a week after installation of the Tiebreaker software. Events that trigger AutoSupport notification include Tiebreaker software panics, detection of disaster conditions on MetroCluster configurations, or an unknown MetroCluster configuration status.

Before you begin

You must have a direct access for setting up AutoSupport messages.

Steps

1. Use the Tiebreaker CLI autosupport command with any of the following options:

Option	Description
-invoke	Sends an AutoSupport message to customer support
-configure wizard	Wizard to configure proxy server credentials
-delete configuration	Deletes the proxy server credentials
--enable	Enables AutoSupport notification (This is the default.)
-disable	Disables AutoSupport notification
-show	Displays AutoSupport status

The following example shows that AutoSupport is enabled or disabled and the destination to which the AutoSupport content is posted:

```
NetApp MetroCluster Tiebreaker :> autosupport enable
AutoSupport already enabled.
```

```
NetApp MetroCluster Tiebreaker :> autosupport disable
AutoSupport status           : disabled
Proxy Server IP Address      : 10.234.168.79
Proxy Server Port Number     : 8090
Proxy Server Username        : admin
AutoSupport destination      :
https://support.netapp.com/asupprod/post/1.0/postAsup
```

```
NetApp MetroCluster Tiebreaker :> autosupport enable
AutoSupport status           : enabled
Proxy Server IP Address      : 10.234.168.79
Proxy Server Port Number     : 8090
Proxy Server Username        : admin
AutoSupport destination      :
https://support.netapp.com/asupprod/post/1.0/postAsup
```

```
NetApp MetroCluster Tiebreaker :> autosupport invoke
AutoSupport transmission     : success
Proxy Server IP Address      : 10.234.168.79
Proxy Server Port Number     : 8090
Proxy Server Username        : admin
AutoSupport destination      :
https://support.netapp.com/asupprod/post/1.0/postAsup
```

The following example shows AutoSupport configured by means of an authenticated proxy server, using an IP address and port number:

```
NetApp MetroCluster Tiebreaker :> autosupport configure wizard
Enter Proxy Server IP address : 10.234.168.79
Enter Proxy Server port number : 8090
Enter Proxy Server Username   : admin
Enter Proxy Server Password   : 123abc
Autosupport configuration updated successfully.
```

The following example shows the deletion of an AutoSupport configuration:

```
NetApp MetroCluster Tiebreaker :> autosupport delete configuration
Autosupport configuration deleted successfully.
```

SNMP events and traps

NetApp MetroCluster Tiebreaker software uses SNMP traps to notify you of significant events. These traps are part of the NetApp MIB file. Each trap contains the following information: trap name, severity, impact level, timestamp, and message.

Event name	Event detail	Trap number
MetroCluster Tie-Breaker is unable to reach the MetroCluster configuration	Warns the administrator that the software cannot detect a disaster. This event occurs when both clusters are not reachable.	25000
MetroCluster Tie-Breaker is unable to reach cluster	Warns the administrator that the software cannot reach one of the clusters.	25001
MetroCluster Tie-Breaker detected disaster at cluster	Notifies the administrator that the software detects a site failure. A notification will be delivered.	25002
All links between partner cluster are severed.	The software detects that both clusters are reachable, but all the network paths between the two clusters are down, and the clusters cannot communicate with each other.	25005
SNMP Test Trap	SNMP configuration can now be tested by running the snmp config test command.	25006

Displaying the status of monitoring operations

You can display the overall status of monitoring operations for a MetroCluster configuration.

Step

1. Use the Tiebreaker CLI monitor show command to display the status of a MetroCluster operation with any of the following options:

Option	Description
-monitor-name	Displays the status for the specified monitor name
-operation-history	Displays up to 10 monitoring operations that were last performed on a cluster
-stats	Displays the statistics related to the specified cluster
-status	Displays the status of the specified cluster Note: The MetroCluster Tiebreaker software might take up to 10 minutes to reflect the completion status of operations such as heal aggregates, heal roots, or switchback.

The following example shows that the clusters cluster_A and cluster_B are connected and healthy:

```
NetApp MetroCluster Tiebreaker:> monitor show -status
MetroCluster: cluster_A
  Disaster: false
  Monitor State: Normal
  Observer Mode: true
  Silent Period: 15
  Override Vetoes: false
  Cluster: cluster_Ba(UUID:4d9ccf24-080f-11e4-9df2-00a098168e7c)
    Reachable: true
    All-Links-Severed: FALSE
      Node: mcc5-a1(UUID:78b44707-0809-11e4-9be1-e50dab9e83e1)
        Reachable: true
        All-Links-Severed: FALSE
        State: normal
      Node: mcc5-a2(UUID:9a8b1059-0809-11e4-9f5e-8d97cdec7102)
        Reachable: true
        All-Links-Severed: FALSE
        State: normal
  Cluster: cluster_B(UUID:70dacd3b-0823-11e4-a7b9-00a0981693c4)
    Reachable: true
    All-Links-Severed: FALSE
      Node: mcc5-b1(UUID:961fce7d-081d-11e4-9ebf-2f295df8fcb3)
        Reachable: true
        All-Links-Severed: FALSE
        State: normal
      Node: mcc5-b2(UUID:9393262d-081d-11e4-80d5-6b30884058dc)
        Reachable: true
        All-Links-Severed: FALSE
        State: normal
```

In the following example, the last seven operations that were run on cluster_B are displayed:

```
NetApp MetroCluster Tiebreaker:> monitor show -operation-history
MetroCluster: cluster_B
 [ 2014-09-15 04:48:32.274 ] MetroCluster Monitor is initialized
 [ 2014-09-15 04:48:32.278 ] Started Discovery and validation of
MetroCluster Setup
 [ 2014-09-15 04:48:35.078 ] Discovery and validation of MetroCluster
Setup succeeded. Started monitoring.
 [ 2014-09-15 04:48:35.246 ] NetApp MetroCluster Tiebreaker software is
able to reach cluster "mcc5a"
 [ 2014-09-15 04:48:35.256 ] NetApp MetroCluster Tiebreaker software is
able to reach cluster "mcc5b"
 [ 2014-09-15 04:48:35.298 ] Link to remote DR cluster is up for cluster
"mcc5a"
 [ 2014-09-15 04:48:35.308 ] Link to remote DR cluster is up for cluster
"mcc5b"
```

Displaying MetroCluster configuration information

You can display the monitor name and IP address of all instances of MetroCluster configurations in the Tiebreaker software.

Step

1. Use the Tiebreaker CLI configuration show command to display the MetroCluster configuration information.

The following example shows the information for clusters cluster_A and cluster_B:

```
MetroCluster: North America
  Monitor Enabled: true
  ClusterA name: cluster_A
  ClusterA IPAddress: 10.222.196.130
  ClusterB name: cluster_B
  ClusterB IPAddress: 10.222.196.140
```

Creating dump files

You save the overall status the Tiebreaker software to a dump file for debugging purposes.

Step

1. Use the Tiebreaker CLI monitor dump -status command to create a dump file of the overall status of all MetroCluster configurations.

The following example shows the successful creation of the /var/log/netapp/mcctb/metrocluster-tiebreaker-status.xml dump file:

```
NetApp MetroCluster Tiebreaker :> monitor dump -status
MetroCluster Tiebreaker status successfully dumped in file
/var/log/netapp/mcctb/metrocluster-tiebreaker-status.xml
```

Disable Tiebreaker observer mode

You can disable observer mode for the Tiebreaker software to automatically initiate a switchover after it detects a site failure.



In Tiebreaker 1.6 and later, you can enable automatic switchover when you add a monitor using the `monitor add wizard` command. Refer to [Add MetroCluster configurations](#).

Step

1. Disable observer mode:

```
monitor modify -monitor-name monitor_name -observer-mode false
```

```
NetApp MetroCluster Tiebreaker :> monitor modify -monitor-name 8pack
-observer-mode false
Warning: If you are turning observer-mode to false, make sure to review
the 'risks and limitations'
as described in the MetroCluster Tiebreaker installation and
configuration.
Are you sure you want to enable automatic switchover capability for
monitor "8pack"? [Y/N]: y
```

Risks and limitations of using MetroCluster Tiebreaker in active mode

Switchover upon detection of a site failure happens automatically, with MetroCluster Tiebreaker in active mode. This mode can be used to supplement the ONTAP/FAS automatic switchover capability.

When you implement MetroCluster Tiebreaker in active mode, the following known issues might lead to data loss:

- When the inter-site link fails, the controllers on each site continue to serve the clients. However, the controllers will not be mirrored. Failure of a controller in one site is identified as a site failure and the MetroCluster Tiebreaker initiates a switchover. The data which is not mirrored after the inter-site link failure with the remote site will be lost.
- A switchover occurs when the aggregates in remote site are in degraded state. The data will not be replicated if the switchover has occurred before aggregate resync.
- A remote storage failure occurs when switchover is in progress.
- The nonvolatile memory (NVRAM or NVMEM, depending on the platform model) in the storage controllers

is not mirrored to the remote disaster recovery (DR) partner on the partner site.

- Metadata is lost if the cluster peering network is down for an extended period and the metadata volumes are not online after a switchover.



You might encounter scenarios that are not mentioned. NetApp is not responsible for any damages that may arise out of use of MetroCluster Tiebreaker in active mode. Do not use MetroCluster Tiebreaker in active mode if the risks and limitations are not acceptable to you.

Firewall requirements for MetroCluster Tiebreaker

MetroCluster Tiebreaker uses a number of ports to communicate with specific services.

The following table lists the ports that you must allow in your firewall:

Port/services	Source	Destination	Purpose
443 / TCP	Tiebreaker	Internet	Sending AutoSupport messages to NetApp
22 / TCP	Management host	Tiebreaker	Tiebreaker Management
443 / TCP	Tiebreaker	Cluster management LIFs	Secure communications to cluster via HTTP (SSL)
22 / TCP	Tiebreaker	Cluster management LIFs	Secure communications to cluster via SSH
443 / TCP	Tiebreaker	Node management LIFs	Secure communications to node via HTTP (SSL)
22 / TCP	Tiebreaker	Node management LIFs	Secure communications to node via SSH
162 / UDP	Tiebreaker	SNMP trap host	Used to send alert notification SNMP traps
ICMP (ping)	Tiebreaker	Cluster management LIFs	Check if cluster IP is reachable
ICMP (ping)	Tiebreaker	Node management LIFs	Check if node IP is reachable

Simulate a switchover using MetroCluster Tiebreaker

Beginning with MetroCluster Tiebreaker 1.7, you can simulate a switchover to test the Tiebreaker switchover functionality.

About this task

- The `monitor switchover-simulate` command is only supported in Tiebreaker 1.7 or later.

Steps

1. Check the Tiebreaker monitoring status:

```
monitor show -status
```

Example

```
NetApp MetroCluster Tiebreaker :> monitor show -status
MetroCluster: A700
  Disaster: false
  Monitor State: Normal
  Observer Mode: false
  Silent Period: 5
  Override Vetoes: false
  Cluster: ClusterA_siteA(UUID:713e5ab2-b4e8-11f0-91aa-00a098ef36a2)
    Reachable: true
    Intersite Connectivity Available: true
      Node: node_A1(UUID:9f6cecbf-b4e4-11f0-9d0f-00a098ef36a2)
        Reachable: true
        Intersite Connectivity Available: true
        State: normal
      Node: node_A2(UUID:2719bb56-b4e7-11f0-996c-00a09897caa3)
        Reachable: true
        Intersite Connectivity Available: true
        State: normal
    Cluster: ClusterB_siteB(UUID:72839591-b4e8-11f0-b688-00a09897cb73)
      Reachable: true
      Intersite Connectivity Available: true
        Node: node_B1(UUID:abfeab89-b4e4-11f0-a077-00a09897cb73)
          Reachable: true
          Intersite Connectivity Available: true
          State: normal
        Node: node_B2(UUID:31e395bf-b4e7-11f0-bf99-00a09897cb2f)
          Reachable: true
          Intersite Connectivity Available: true
          State: normal
```

2. Perform a switchover simulation:

```
monitor switchover-simulate -monitor-name <monitor_name> -cluster
<cluster_name>
```

The command returns the following output for a successful operation:

```
Successfully triggered Switchover Simulation. Please check the status of
the Switchover Simulation on the ONTAP cluster using command
"metrocluster operation history show"
```

Event log files for MetroCluster Tiebreaker

The event log file contains a log of all the actions performed by the MetroCluster Tiebreaker software.

The Tiebreaker software performs the following actions:

- Detects site disasters
- Detects configuration changes related to the database, other Tiebreaker monitors, or the MetroCluster Tiebreaker software
- Detects SSH connections and disconnections
- Discovers MetroCluster configurations

These actions are logged in the event log file in the following format:

```
<timestamp> <severity/log level> <thread-id> <module>
```

Example

```
2022-09-07 06:14:30,797 INFO [MCCTBCommandServer-16] [SslSupport]
Successfully initiated SSL context. Protocol used is TLSv1.3.
2022-09-07 06:14:34,137 INFO [MCCTBCommandServer-16] [DataBase]
Successfully read MCCTB database.
2022-09-07 06:14:34,137 INFO [MCCTBCommandServer-16]
[ConfigurationMonitor] Debug mode disabled.
```

Where to find additional information

You can learn more about MetroCluster configuration and operation.

MetroCluster and miscellaneous information

Information	Subject
MetroCluster Documentation	<ul style="list-style-type: none">• All MetroCluster information
NetApp Technical Report 4375: NetApp MetroCluster for ONTAP 9.3	<ul style="list-style-type: none">• A technical overview of the MetroCluster configuration and operation.• Best practices for MetroCluster configuration.

Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the MetroCluster IP configuration • Configuring the MetroCluster in ONTAP
Maintain the MetroCluster components	<ul style="list-style-type: none"> • Guidelines for maintenance in a MetroCluster configuration • Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches • Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster configuration • Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster configuration • Replacing hardware at a disaster site in a fabric-attached or stretch MetroCluster configuration • Expanding a two-node fabric-attached or stretch MetroCluster configuration to a four-node MetroCluster configuration. • Expanding a four-node fabric-attached or stretch MetroCluster configuration to an eight-node MetroCluster configuration.
Active IQ Unified Manager documentation NetApp Documentation: Product Guides and Resources	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems

Understand MetroCluster data protection and disaster recovery

Understanding MetroCluster data protection and disaster recovery

It is helpful to understand how MetroCluster protects data and provides transparent recovery from failures so that you can manage your switchover and switchback activities easily and efficiently.

MetroCluster uses mirroring to protect the data in a cluster. It provides disaster recovery through a single MetroCluster command that activates a secondary on the survivor site to serve the mirrored data originally owned by a primary site affected by disaster.

How eight- and four-node MetroCluster configurations provide local failover and switchover

Eight- and four-node MetroCluster configurations protect data on both a local level and cluster level. If you are setting up a MetroCluster configuration, you need to know how MetroCluster configurations protect your data.

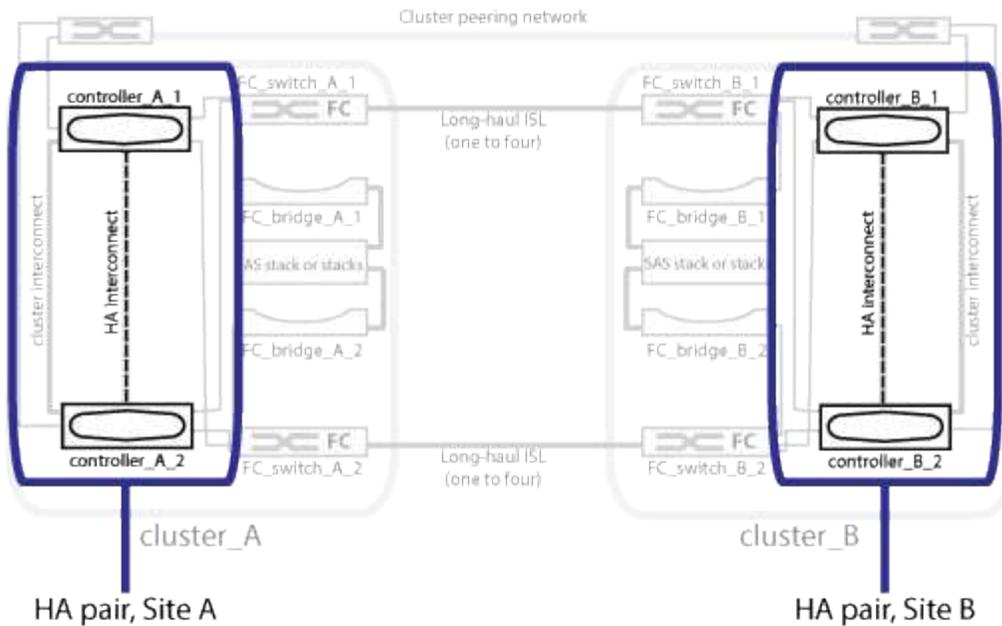
MetroCluster configurations protect data by using two physically separated, mirrored clusters. Each cluster synchronously mirrors the data and storage virtual machine (SVM) configuration of the other. When a disaster occurs at one site, an administrator can activate the mirrored SVM and begin serving the mirrored data from the surviving site. Additionally, the nodes in each cluster are configured as an HA pair, providing a level of local failover.

How local HA data protection works in a MetroCluster configuration

You need to understand how HA pairs work in the MetroCluster configuration.

The two clusters in the peered network provide bidirectional disaster recovery, where each cluster can be the source and backup of the other cluster. Each cluster includes two nodes, which are configured as an HA pair. In the case of a failure or required maintenance within a single node's configuration, storage failover can transfer that node's operations to its local HA partner.

The following illustration shows a MetroCluster FC configuration. The HA functionality is the same in MetroCluster IP configurations, except that the HA interconnect is provided by the cluster switches.



Related information

[High-availability configuration](#)

How MetroCluster configurations provide data and configuration replication

MetroCluster configurations use a variety of ONTAP features to provide synchronous replication of data and configuration between the two MetroCluster sites.

Configuration protection with the configuration replication service

The ONTAP configuration replication service (CRS) protects the MetroCluster configuration by automatically replicating the information to the DR partner.

The CRS synchronously replicates local node configuration to the DR partner in the partner cluster. This replication is carried out over the cluster peering network.

The information replicated includes the cluster configuration and the SVM configuration.

Replication of SVMs during MetroCluster operations

The ONTAP configuration replication service (CRS) provides redundant data server configuration and mirroring of data volumes that belong to the SVM. If a switchover occurs, the source SVM is brought down and the destination SVM, located on the surviving cluster, becomes active.



Destination SVMs in the MetroCluster configuration have the suffix “-mc” automatically appended to their name to help identify them. A MetroCluster configuration appends the suffix “-mc” to the name of the destination SVMs, if the SVM name contains a period, the suffix “-mc” is applied prior to the first period. For example, if the SVM name is SVM.DNS.NAME, then the suffix “-mc” is appended as SVM-MC.DNS.NAME.

The following example shows the SVMs for a MetroCluster configuration, where “SVM_cluster_A” is an SVM on the source site and “SVM_cluster_A-mc” is a sync-destination aggregate on the disaster recovery site.

- SVM_cluster_A serves data on cluster A.

It is a sync-source SVM that represents the SVM configuration (LIFs, protocols, and services) and data in volumes belonging to the SVM. The configuration and data are replicated to SVM_cluster_A-mc, a sync-destination SVM located on cluster B.

- SVM_cluster_B serves data on cluster B.

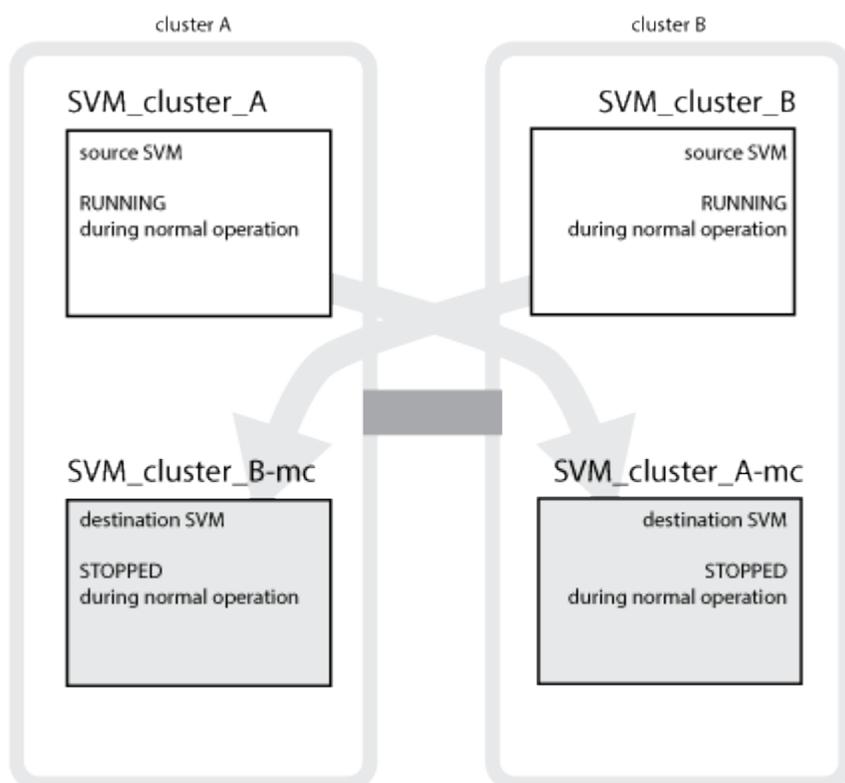
It is a sync-source SVM that represents configuration and data to SVM_cluster_B-mc located on cluster A.

- SVM_cluster_B-mc is a sync-destination SVM that is stopped during normal, healthy operation of the MetroCluster configuration.

In a successful switchover from cluster B to cluster A, SVM_cluster_B is stopped and SVM_cluster_B-mc is activated and begins serving data from cluster A.

- SVM_cluster_A-mc is a sync-destination SVM that is stopped during normal, healthy operation of the MetroCluster configuration.

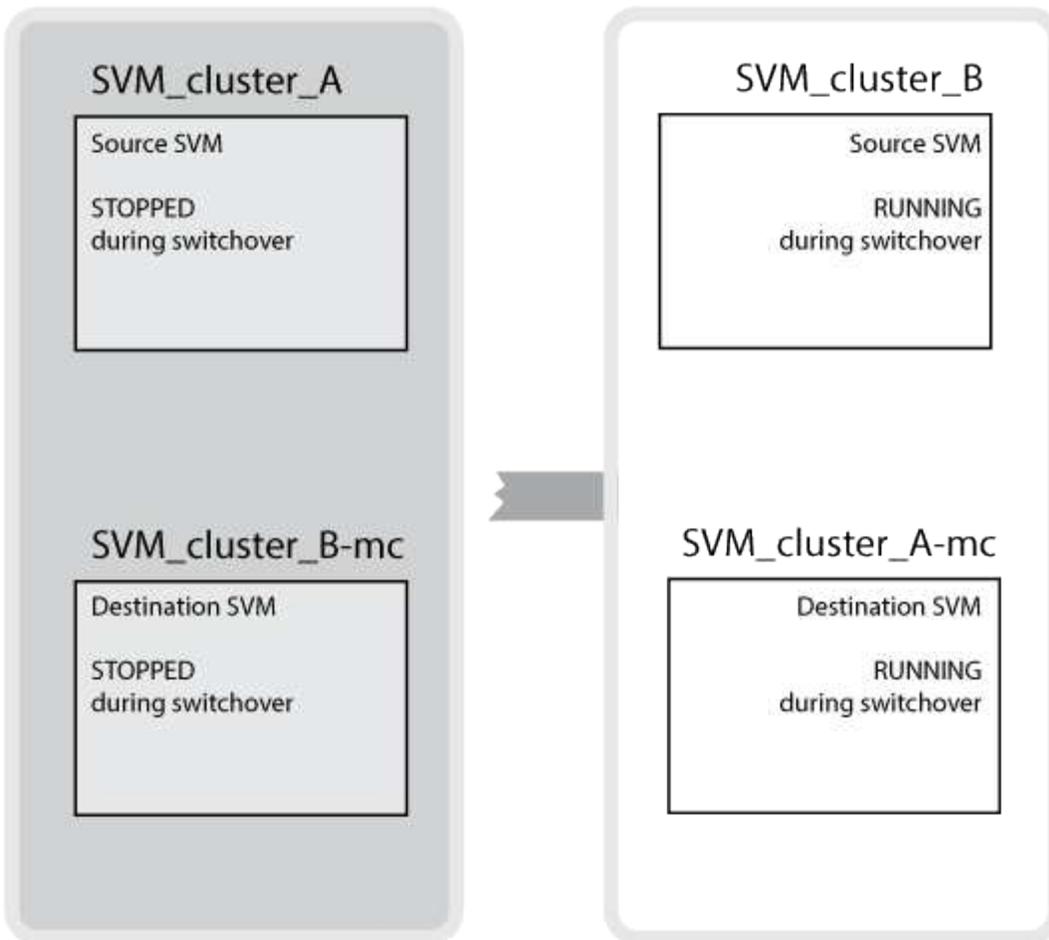
In a successful switchover from cluster A to cluster B, SVM_cluster_A is stopped and SVM_cluster_A-mc is activated and begins serving data from cluster B.



If a switchover occurs, the remote plex on the surviving cluster comes online and the secondary SVM begins serving the data.

cluster A DOWN AND SWITCHED OVER

cluster B UP



The availability of remote plexes after switchover depends on the MetroCluster configuration type:

- For MetroCluster FC configurations, after switchover, both local and remote plexes remain online if the disaster site storage is accessible via the ISLs.

If the ISLs have failed and the disaster site storage is not available, the sync-destination SVM begins serving data from the surviving site.

- For MetroCluster IP configurations the availability of the remote plexes depends on the ONTAP version:
 - Beginning with ONTAP 9.5, both local and remote plexes remain online if the disaster site nodes remain booted up.
 - Prior to ONTAP 9.5, storage is available only from local plex on the surviving site.

The sync-destination SVM begins serving data from the surviving site.

Related information

[System administration](#)

How MetroCluster configurations use SyncMirror to provide data redundancy

Mirrored aggregates using SyncMirror functionality provide data redundancy and contain the volumes owned by the source and destination storage virtual machine (SVM). Data is replicated into disk pools on the partner cluster. Unmirrored aggregates are also supported.

The following table shows the state (online or offline) of an unmirrored aggregate after a switchover:

Type of switchover	MetroCluster FC configuration state	MetroCluster IP configuration state
Negotiated switchover (NSO)	Online	Offline (Note 1)
Automatic unplanned switchover (AUSO)	Online	Offline (Note 1)
Unplanned switchover (USO)	<ul style="list-style-type: none"> If storage is not available: Offline If storage is available: Online 	Offline (Note 1)

Note 1: In MetroCluster IP configurations, after the switchover is complete, you can manually bring the unmirrored aggregates online.

Learn more about [Differences in switchover between MetroCluster FC and IP configurations](#).

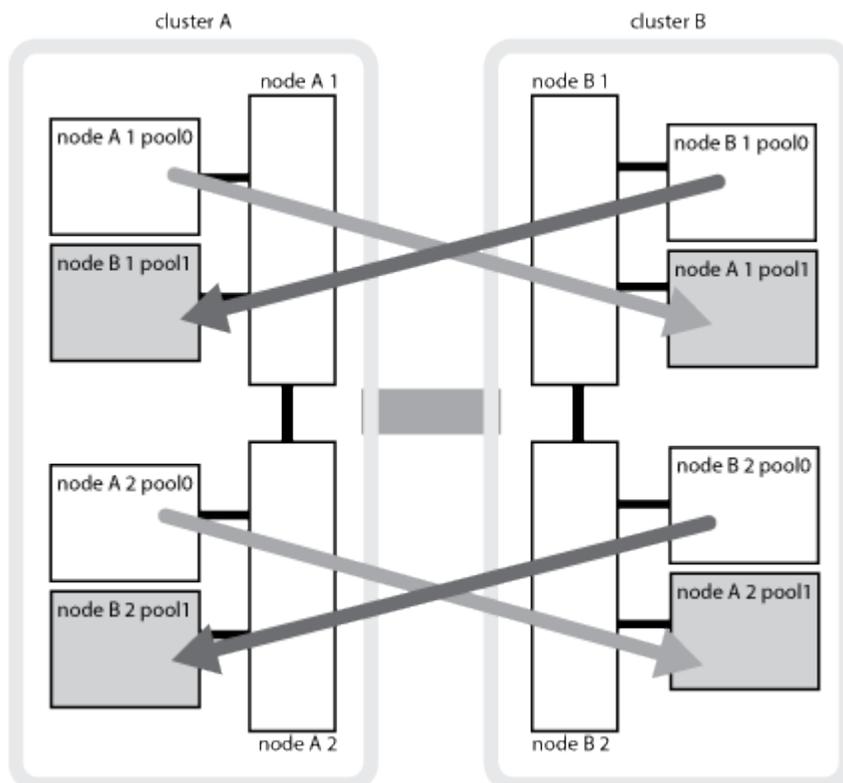


After a switchover, if the unmirrored aggregate is at the DR partner node and there is an inter-switch link (ISL) failure, then that local node might fail.

The following illustration shows how disk pools are mirrored between the partner clusters. Data in local plexes (in pool0) is replicated to remote plexes (in pool1).



If hybrid aggregates are used, performance degradation can occur after a SyncMirror plex has failed due to the solid-state disk (SSD) layer filling up.

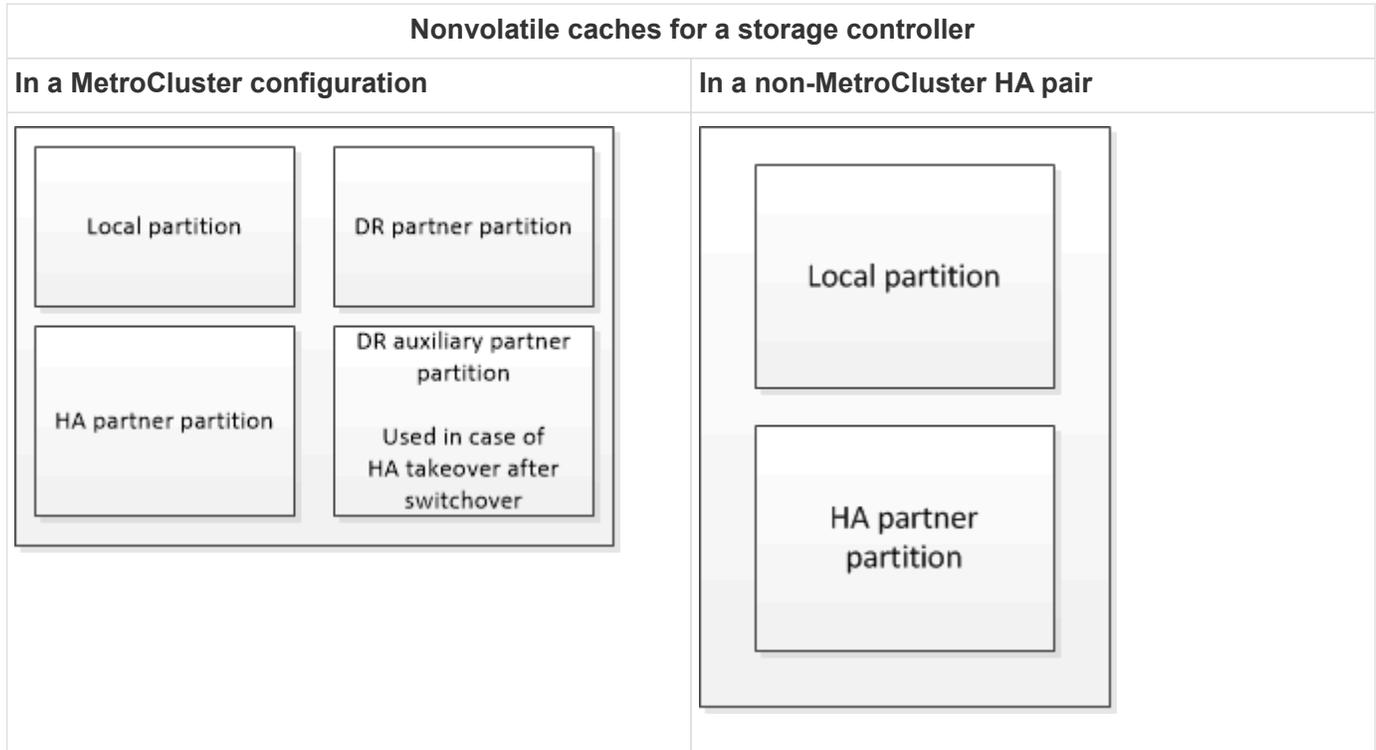


How NVRAM or NVMEM cache mirroring and dynamic mirroring work in MetroCluster configurations

The nonvolatile memory (NVRAM or NVMEM, depending on the platform model) in the storage controllers is mirrored both locally to a local HA partner and remotely to a remote disaster recovery (DR) partner on the partner site. In the event of a local failover or switchover, this configuration enables data in the nonvolatile cache to be preserved.

In an HA pair that is not part of a MetroCluster configuration, each storage controller maintains two nonvolatile cache partitions: one for itself and one for its HA partner.

In a four-node MetroCluster configuration, the nonvolatile cache of each storage controller is divided into four partitions. In a two-node MetroCluster configuration, the HA partner partition and DR auxiliary partition are not used, because the storage controllers are not configured as an HA pair.



The nonvolatile caches store the following content:

- The local partition holds data that the storage controller has not yet written to disk.
- The HA partner partition holds a copy of the local cache of the storage controller's HA partner.

In a two-node MetroCluster configuration, there is no HA partner partition because the storage controllers are not configured as an HA pair.

- The DR partner partition holds a copy of the local cache of the storage controller's DR partner.

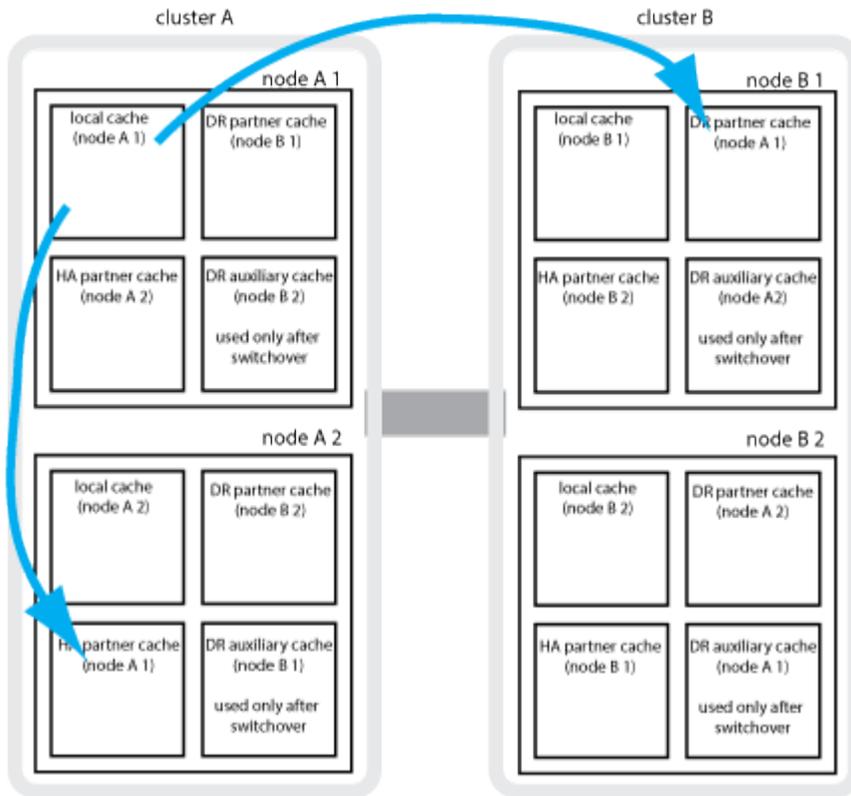
The DR partner is a node in the partner cluster that is paired with the local node.

- The DR auxiliary partner partition holds a copy of the local cache of the storage controller's DR auxiliary partner.

The DR auxiliary partner is the HA partner of the local node's DR partner. This cache is needed if there is an HA takeover (either when the configuration is in normal operation or after a MetroCluster switchover).

In a two-node MetroCluster configuration, there is no DR auxiliary partner partition because the storage controllers are not configured as an HA pair.

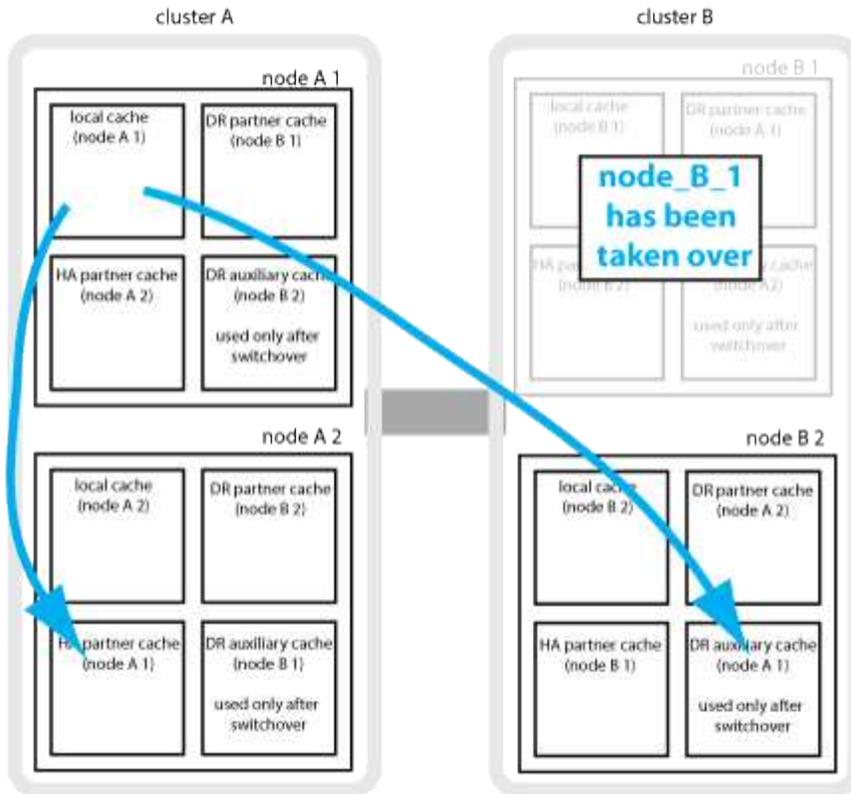
For example, the local cache of a node (node_A_1) is mirrored both locally and remotely at the MetroCluster sites. The following illustration shows that the local cache of node_A_1 is mirrored to the HA partner (node_A_2) and DR partner (node_B_1):



Dynamic mirroring in event of a local HA takeover

If a local HA takeover occurs in a four-node MetroCluster configuration, the taken-over node can no longer act as a mirror for its DR partner. To allow DR mirroring to continue, the mirroring automatically switches to the DR auxiliary partner. After a successful giveback, mirroring automatically returns to the DR partner.

For example, node_B_1 fails and is taken over by node_B_2. The local cache of node_A_1 can no longer be mirrored to node_B_1. The mirroring switches to the DR auxiliary partner, node_B_2.



Types of disasters and recovery methods

You need to be familiar with different types of failures and disasters so that you can use the MetroCluster configuration to respond appropriately.

- Single-node failure

A single component in the local HA pair fails.

In a four-node MetroCluster configuration, this failure might lead to an automatic or a negotiated takeover of the impaired node, depending on the component that failed. Data recovery is described in [High-availability pair management](#).

In a two-node MetroCluster configuration, this failure leads to an automatic unplanned switchover (AUSO).

- Site-wide controller failure

All controller modules fail at a site because of loss of power, replacement of equipment, or disaster. Typically, MetroCluster configurations cannot differentiate between failures and disasters. However, witness software, such as the MetroCluster Tiebreaker software, can differentiate between them. A site-wide controller failure condition can lead to an automatic switchover if Inter-Switch Link (ISL) links and switches are up and the storage is accessible.

[High-availability pair management](#) has more information about how to recover from site-wide controller failures that do not include controller failures, as well as failures that include one or more controllers.

- ISL failure

The links between the sites fail. The MetroCluster configuration takes no action. Each node continues to serve data normally, but the mirrors are not written to the respective disaster recovery sites because

access to them is lost.

- Multiple sequential failures

Multiple components fail in a sequence. For example, a controller module, a switch fabric, and a shelf fail in a sequence and result in a storage failover, fabric redundancy, and SyncMirror sequentially protecting against downtime and data loss.

The following table shows failure types, and the corresponding disaster recovery (DR) mechanism and recovery method:



AUSO (automatic unplanned switchover) is not supported on MetroCluster IP configurations.

Failure type	DR mechanism		Summary of recovery method	
	Four-node configuration	Two-node configuration	Four-node configuration	Two-node configuration
Single-node failure	Local HA failover	AUSO	Not required if automatic failover and giveback is enabled.	<p>After the node is restored, manual healing and switchback using the <code>metrocluster heal -phase aggregates</code>, <code>metrocluster heal -phase root-aggregates</code>, and <code>metrocluster switchback</code> commands is required.</p> <p>NOTE: The <code>metrocluster heal</code> commands are not required on MetroCluster IP configurations running ONTAP 9.5 or later.</p>

Site failure	MetroCluster switchover		After the node is restored, manual healing and switchback using the <code>metrocluster healing</code> and <code>metrocluster switchback</code> commands is required. The <code>metrocluster heal</code> commands are not required on MetroCluster IP configurations running ONTAP 9.5.
Site-wide controller failure	AUSO Only if the storage at the disaster site is accessible.	AUSO (same as single-node failure)	
Multiple sequential failures	Local HA failover followed by MetroCluster forced switchover using the <code>metrocluster switchover -forced -on-disaster</code> command. NOTE: Depending on the component that failed, a forced switchover might not be required.	MetroCluster forced switchover using the <code>metrocluster switchover -forced-on -disaster</code> command.	
ISL failure	No MetroCluster switchover; the two clusters independently serve their data		

How an eight-node or four-node MetroCluster configuration provides nondisruptive operations

In the case of an issue limited to a single node, a failover and giveback within the local HA pair provides continued nondisruptive operation. In this case, the MetroCluster configuration does not require a switchover to the remote site.

Because the eight-node or four-node MetroCluster configuration consists of one or more HA pair at each site, each site can withstand local failures and perform nondisruptive operations without requiring a switchover to the partner site. The operation of the HA pair is the same as HA pairs in non-MetroCluster configurations.

For four-node and eight-node MetroCluster configurations, node failures due to panic or power loss can cause an automatic switchover.

High-availability pair management

If a second failure occurs after a local failover, the MetroCluster switchover event provides continued nondisruptive operations. Similarly, after a switchover operation, in the event of a second failure in one of the surviving nodes, a local failover event provides continued nondisruptive operations. In this case, the single surviving node serves data for the other three nodes in the DR group.

Switchover and switchback during MetroCluster transition

MetroCluster FC-to-IP transition involves adding MetroCluster IP nodes and IP switches to an existing MetroCluster FC configuration, and then retiring the MetroCluster FC nodes. Depending on the stage of the transition process, the MetroCluster switchover, healing, and switchback operations use different workflows.

See [Switchover, healing, and switchback operations during transition](#).

Consequences of local failover after switchover

If a MetroCluster switchover occurs, and then an issue arises at the surviving site, a local failover can provide continued, nondisruptive operation. However, the system is at risk because it is no longer in a redundant configuration.

If a local failover occurs after a switchover has occurred, a single controller serves data for all storage systems in the MetroCluster configuration, leading to possible resource issues, and is vulnerable to additional failures.

How a two-node MetroCluster configuration provides nondisruptive operations

If one of the two sites has an issue due to panic, the MetroCluster switchover provides continued nondisruptive operation. If the power loss impacts both the node and the storage, then the switchover is not automatic and there is a disruption until the `metrocluster switchover` command is issued.

Because all storage is mirrored, a switchover operation can be used to provide nondisruptive resiliency in case of a site failure similar to that found in a storage failover in an HA pair for a node failure.

For two-node configurations, the same events that trigger an automatic storage failover in an HA pair trigger an automatic unplanned switchover (AUSO). This means that a two-node MetroCluster configuration has the same level of protection as an HA pair.

Related information

[Automatic unplanned switchover in MetroCluster FC configurations](#)

Overview of the switchover process

The MetroCluster switchover operation enables immediate resumption of services following a disaster by moving storage and client access from the source cluster to the remote site. You must be aware of what changes to expect and which actions you need to perform if a switchover occurs.

During a switchover operation, the system takes the following actions:

- Ownership of the disks that belong to the disaster site is changed to the disaster recovery (DR) partner.

This is similar to the case of a local failover in a high-availability (HA) pair, in which ownership of the disks belonging to the partner that is down is changed to the healthy partner.

- The surviving plexes that are located on the surviving site but belong to the nodes in the disaster cluster are brought online on the cluster at the surviving site.
- The sync-source storage virtual machine (SVM) that belongs to the disaster site is brought down only during a negotiated switchover.



This is applicable only to a negotiated switchover.

- The sync-destination SVM belonging to the disaster site is brought up.

While being switched over, the root aggregates of the DR partner are not brought online.

The `metrocluster switchover` command switches over the nodes in all DR groups in the MetroCluster configuration. For example, in an eight-node MetroCluster configuration, it switches over the nodes in both DR groups.

If you are switching over only services to the remote site, you should perform a negotiated switchover without

fencing the site. If storage or equipment is unreliable, you should fence the disaster site, and then perform an unplanned switchover. Fencing prevents RAID reconstructions when the disks power up in a staggered manner.



This procedure should be only used if the other site is stable and not intended to be taken offline.

Availability of commands during switchover

The following table shows the availability of commands during switchover:

Command	Availability
<code>storage aggregate create</code>	<p>You can create an aggregate:</p> <ul style="list-style-type: none"> • If it is owned by a node that is part of the surviving cluster <p>You cannot create an aggregate:</p> <ul style="list-style-type: none"> • For a node at the disaster site • For a node that is part of the surviving cluster
<code>storage aggregate delete</code>	You can delete a data aggregate.
<code>storage aggregate mirror</code>	You can create a plex for a non-mirrored aggregate.
<code>storage aggregate plex delete</code>	You can delete a plex for a mirrored aggregate.
<code>vserver create</code>	<p>You can create an SVM:</p> <ul style="list-style-type: none"> • If its root volume resides in a data aggregate owned by the surviving cluster <p>You cannot create an SVM:</p> <ul style="list-style-type: none"> • If its root volume resides in a data aggregate owned by the disaster-site cluster
<code>vserver delete</code>	You can delete both sync-source and sync-destination SVMs.
<code>network interface create -lif</code>	You can create a data SVM LIF for both sync-source and sync-destination SVMs.
<code>network interface delete -lif</code>	You can delete a data SVM LIF for both sync-source and sync-destination SVMs.

volume create	<p>You can create a volume for both sync-source and sync-destination SVMs.</p> <ul style="list-style-type: none"> • For a sync-source SVM, the volume must reside in a data aggregate owned by the surviving cluster • For a sync-destination SVM, the volume must reside in a data aggregate owned by the disaster-site cluster
volume delete	<p>You can delete a volume for both sync-source and sync-destination SVMs.</p>
volume move	<p>You can move a volume for both sync-source and sync-destination SVMs.</p> <ul style="list-style-type: none"> • For a sync-source SVM, the surviving cluster must own the destination aggregate • For a sync-destination SVM, the disaster-site cluster must own the destination aggregate
snapmirror break	<p>You can break a SnapMirror relationship between a source and destination endpoint of a data protection mirror.</p>

Differences in switchover between MetroCluster FC and IP configurations

In MetroCluster IP configurations, because the remote disks are accessed through the remote DR partner nodes acting as iSCSI targets, the remote disks are not accessible when the remote nodes are taken down in a switchover operation. This results in differences with MetroCluster FC configurations:

- Mirrored aggregates that are owned by the local cluster become degraded.
- Mirrored aggregates that were switched over from the remote cluster become degraded.



When unmirrored aggregates are supported on a MetroCluster IP configuration, the unmirrored aggregates that are not switched over from the remote cluster are not accessible.

Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration

The ownership of disks temporarily changes automatically during high availability and MetroCluster operations. It is helpful to know how the system tracks which node owns which disks.

In ONTAP, a controller module's unique system ID (obtained from a node's NVRAM card or NVMEM board) is used to identify which node owns a specific disk. Depending on the HA or DR state of the system, the ownership of the disk might temporarily change. If the ownership changes because of an HA takeover or a DR switchover, the system records which node is the original (called "home") owner of the disk, so that it can return the ownership after HA giveback or DR switchback. The system uses the following fields to track disk ownership:

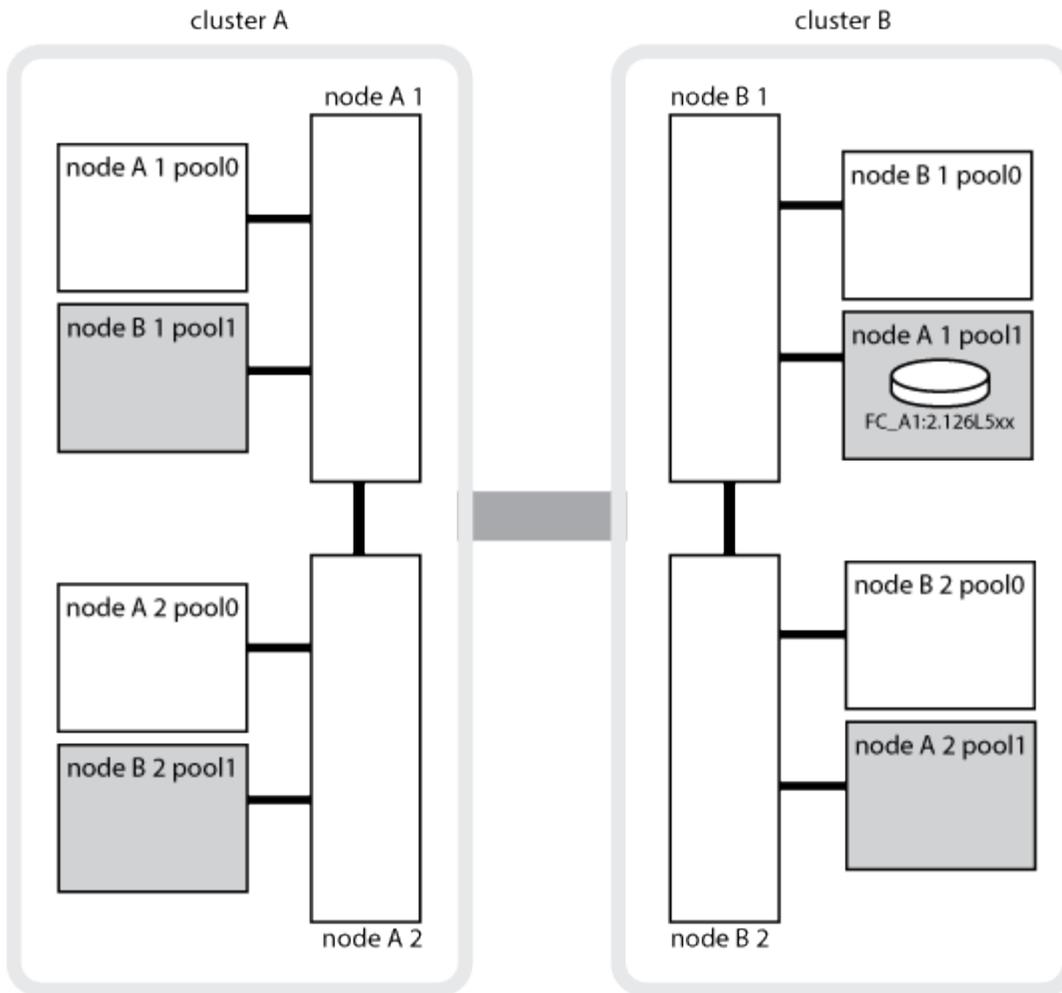
- Owner
- Home owner
- DR Home owner

In the MetroCluster configuration, in the event of a switchover, a node can take ownership of an aggregate originally owned by nodes in the partner cluster. Such aggregates are referred to as cluster-foreign aggregates. The distinguishing feature of a cluster-foreign aggregate is that it is an aggregate not currently known to the cluster, and so the DR Home owner field is used to show that it is owned by a node from the partner cluster. A traditional foreign aggregate within an HA pair is identified by Owner and Home owner values being different, but the Owner and Home owner values are the same for a cluster-foreign aggregate; thus, you can identify a cluster-foreign aggregate by the DR Home owner value.

As the state of the system changes, the values of the fields change, as shown in the following table:

Field	Value during...			
	Normal operation	Local HA takeover	MetroCluster switchover	Takeover during switchover
Owner	ID of the node that has access to the disk.	ID of the HA partner, which temporarily has access to the disk.	ID of the DR partner, which temporarily has access to the disk.	ID of the DR auxiliary partner, which temporarily has access to the disk.
Home owner	ID of the original owner of the disk within the HA pair.	ID of the original owner of the disk within the HA pair.	ID of the DR partner, which is the Home owner in the HA pair during the switchover.	ID of the DR partner, which is the Home owner in the HA pair during the switchover.
DR Home owner	Empty	Empty	ID of the original owner of the disk within the MetroCluster configuration.	ID of the original owner of the disk within the MetroCluster configuration.

The following illustration and table provide an example of how ownership changes, for a disk in node_A_1's disk pool1, physically located in cluster_B.



MetroCluster state	Owner	Home owner	DR Home owner	Notes
Normal with all nodes fully operational.	node_A_1	node_A_1	not applicable	
Local HA takeover, node_A_2 has taken over disks belonging to its HA partner node_A_1.	node_A_2	node_A_1	not applicable	
DR switchover, node_B_1 has taken over disks belong to its DR partner, node_A_1.	node_B_1	node_B_1	node_A_1	The original home node ID is moved to the DR Home owner field. After aggregate switchback or healing, ownership goes back to node_A_1.

In DR switchover and local HA takeover (double failure), node_B_2 has taken over disks belonging to its HA node_B_1.	node_B_2	node_B_1	node_A_1	After giveback, ownership goes back to node_B_1. After switchback or healing, ownership goes back to node_A_1.
After HA giveback and DR switchback, all nodes fully operational.	node_A_1	node_A_1	not applicable	

Considerations when using unmirrored aggregates

If your configuration includes unmirrored aggregates, you must be aware of potential access issues after switchover operations.

Considerations for unmirrored aggregates when doing maintenance requiring power shutdown

If you are performing negotiated switchover for maintenance reasons requiring site-wide power shutdown, you should first manually take offline any unmirrored aggregates owned by the disaster site.

If you do not, nodes at the surviving site might go down due to multi-disk panics. This could occur if switched-over unmirrored aggregates go offline or are missing because of the loss of connectivity to storage at the disaster site due to the power shutdown or a loss of ISLs.

Considerations for unmirrored aggregates and hierarchical namespaces

If you are using hierarchical namespaces, you should configure the junction path so that all of the volumes in that path are either on mirrored aggregates only or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates in the junction path might prevent access to the unmirrored aggregates after the switchover operation.

Considerations for unmirrored aggregates and CRS metadata volume and data SVM root volumes

The configuration replication service (CRS) metadata volume and data SVM root volumes must be on a mirrored aggregate. You cannot move these volumes to unmirrored aggregate. If they are on unmirrored aggregate, negotiated switchover and switchback operations are vetoed. The `metrocluster check` command provides a warning if this is the case.

Considerations for unmirrored aggregates and SVMs

SVMs should be configured on mirrored aggregates only or on unmirrored aggregates only. Configuring a mix of unmirrored and mirrored aggregates can result in a switchover operation that exceeds 120 seconds and result in a data outage if the unmirrored aggregates do not come online.

Considerations for unmirrored aggregates and SAN

A LUN should not be located on an unmirrored aggregate. Configuring a LUN on an unmirrored aggregate can result in a switchover operation that exceeds 120 seconds and a data outage.

Automatic unplanned switchover in MetroCluster FC configurations

In MetroCluster FC configurations, certain scenarios can trigger an automatic unplanned switchover (AUSO) in the event of a site-wide controller failure to provide nondisruptive operations. AUSO can be disabled if desired.



Automatic unplanned switchover is not supported in MetroCluster IP configurations.

In a MetroCluster FC configuration, an AUSO can be triggered if all nodes at a site are failed because of the following reasons:

- Power down
- Power loss
- Panic



In an eight-node MetroCluster FC configuration, you can set an option to trigger an AUSO if both nodes in an HA pair fail.

Because there is no local HA failover available in a two-node MetroCluster configuration, the system performs an AUSO to provide continued operation after a controller failure. This functionality is similar to the HA takeover capability in an HA pair. In a two-node MetroCluster configuration, an AUSO can be triggered in the following scenarios:

- Node power down
- Node power loss
- Node panic
- Node reboot

If an AUSO occurs, disk ownership for the impaired node's pool0 and pool1 disks is changed to the disaster recovery (DR) partner. This ownership change prevents the aggregates from going into a degraded state after the switchover.

After the automatic switchover, you must manually proceed through the healing and switchback operations to return the controller to normal operation.

Hardware-assisted AUSO in two-node MetroCluster configurations

In a two-node MetroCluster configuration, the controller module's service processor (SP) monitors the configuration. In some scenarios, the SP can detect a failure faster than the ONTAP software. In this case, the SP triggers AUSO. This feature is automatically enabled.

The SP sends and receives SNMP traffic to and from its DR partner to monitor its health.

Changing the AUSO setting in MetroCluster FC configurations

AUSO is set to "auso-on-cluster-disaster" by default. Its status can be viewed in the `metrocluster show` command.



The AUSO setting does not apply to MetroCluster IP configurations.

You can disable AUSO with the `metrocluster modify -auto-switchover-failure-domain auto-disabled` command. This command prevents triggering AUSO in DR site-wide controller failure. It should be

run on both the sites if you want to disable AUSO on both the sites.

AUSO can be reenabled with the `metrocluster modify -auto-switchover-failure-domain auso-on-cluster-disaster` command.

AUSO can also be set to “auso-on-dr-group-disaster”. This advance level command triggers AUSO on HA failover at one site. It should be run on both the sites with the `metrocluster modify -auto-switchover-failure-domain auso-on-dr-group-disaster` command.

The AUSO setting during switchover

When switchover occurs, the AUSO setting is disabled internally because if a site is in switchover, it cannot automatically switch over.

Recovering from AUSO

To recover from an AUSO, you perform the same steps as for a planned switchover.

[Performing switchover for tests or maintenance](#)

Mediator-assisted automatic unplanned switchover in MetroCluster IP configurations

[Learn about how the ONTAP Mediator supports automatic unplanned switchover in MetroCluster IP configurations.](#)

What happens during healing (MetroCluster FC configurations)

During healing in MetroCluster FC configurations, the resynchronization of mirrored aggregates occurs in a phased process that prepares the nodes at the repaired disaster site for switchback. It is a planned event, thereby giving you full control of each step to minimize downtime. Healing is a two-step process that occurs on the storage and controller components.

Data aggregate healing

After the problem at the disaster site is resolved, you start the storage healing phase:

1. Checks that all nodes are up and running at the surviving site.
2. Changes ownership of all the pool 0 disks at the disaster site, including root aggregates.

During this phase of healing, the RAID subsystem resynchronizes mirrored aggregates, and the WAFL subsystem replays the nvsave files of mirrored aggregates that had a failed pool 1 plex at the time of switchover.

If some source storage components failed, the command reports the errors at applicable levels: Storage, Sanown, or RAID.

If no errors are reported, the aggregates are successfully resynchronized. This process can sometimes take hours to complete.

[Healing the configuration](#)

Root aggregate healing

After the aggregates are synchronized, you start the controller healing phase by giving back the CFO

aggregates and root aggregates to their respective DR partners.

Healing the configuration

What happens during healing (MetroCluster IP configurations)

During healing in MetroCluster IP configurations, the resynchronization of mirrored aggregates occurs in a phased process that prepares the nodes at the repaired disaster site for switchback. It is a planned event, thereby giving you full control of each step to minimize downtime. Healing is a two-step process that occurs on the storage and controller components.

Differences with MetroCluster FC configurations

In MetroCluster IP configurations, you must boot the nodes in the disaster site cluster before the healing operation is performed.

The nodes in the disaster site cluster must be running so that the remote iSCSI disks can be accessed when aggregates are resynchronized.

If the disaster site nodes are not running, the healing operation fails because the disaster node cannot perform the disk ownership changes needed.

Data aggregate healing

After the problem at the disaster site is resolved, you start the storage healing phase:

1. Checks that all nodes are up and running at the surviving site.
2. Changes ownership of all the pool 0 disks at the disaster site, including root aggregates.

During this phase of healing, the RAID subsystem resynchronizes mirrored aggregates, and the WAFL subsystem replays the nvsave files of mirrored aggregates that had a failed pool 1 plex at the time of switchover.

If some source storage components failed, the command reports the errors at applicable levels: Storage, Sanown, or RAID.

If no errors are reported, the aggregates are successfully resynchronized. This process can sometimes take hours to complete.

Healing the configuration

Root aggregate healing

After the aggregates are synchronized, you perform the root aggregate healing phase. In MetroCluster IP configurations, this phase confirms that aggregates have been healed.

Healing the configuration

Automatic healing of aggregates on MetroCluster IP configurations after switchover

Beginning with ONTAP 9.5, healing is automated during negotiated switchover operations on MetroCluster IP configurations. Beginning with ONTAP 9.6, automated healing after unscheduled switchover is supported. This removes the requirement to issue the `metrocluster heal` commands.

Automatic healing after negotiated switchover (beginning with ONTAP 9.5)

After performing a negotiated switchover (a switchover command issued without the `-forced-on-disaster true` option), the automatic healing functionality simplifies the steps required to return the system to normal operation. On systems with automatic healing, the following occurs after the switchover:

- The disaster site nodes remain up.

Because they are in switchover state, they are not serving data from their local mirrored plexes.

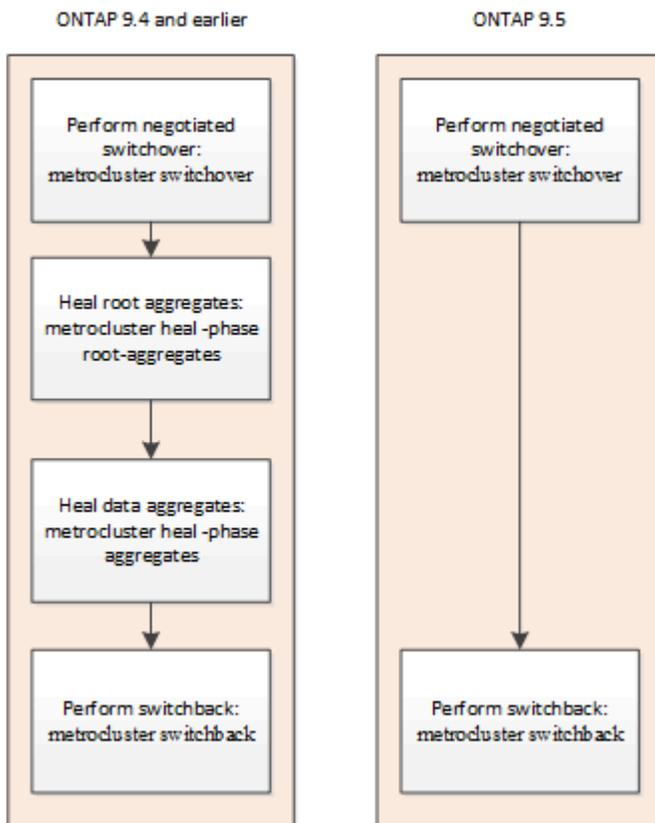
- The disaster site nodes are moved to the “Waiting for switchback” state.

You can confirm the status of the disaster site nodes by using the `metrocluster operation show` command.

- You can perform the switchback operation without issuing the healing commands.

This feature applies to MetroCluster IP configurations running ONTAP 9.5 and later. It does not apply to MetroCluster FC configurations.

The manual healing commands are still required on MetroCluster IP configurations running ONTAP 9.4 and earlier.



Automatic healing after unscheduled switchover (beginning with ONTAP 9.6)

Automatic healing after an unscheduled switchover is supported on MetroCluster IP configurations beginning with ONTAP 9.6. An unscheduled switchover is one in which you issue the `switchover` command with the `-forced-on-disaster true` option.

Automatic healing after an unscheduled switchover is not supported on MetroCluster FC configurations, and the manual healing commands are still required after unscheduled switchover on MetroCluster IP

configurations running ONTAP 9.5 and earlier.

On systems running ONTAP 9.6 and later, the following occurs after the unscheduled switchover:

- Depending on the extent of the disaster, the disaster site nodes can be down.

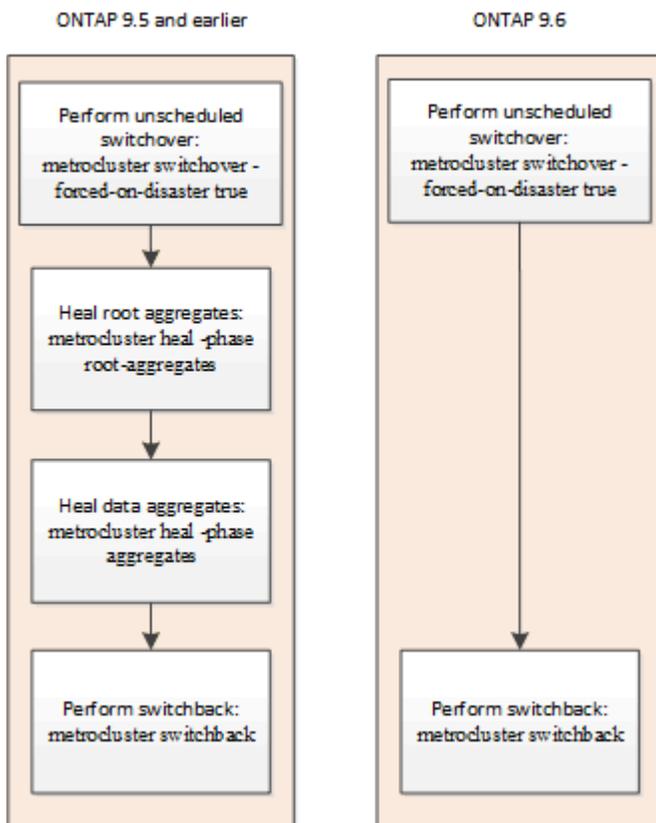
Because they are in switchover state, they are not serving data from their local mirrored plexes, even if they are powered up.

- If the disaster sites were down, when booted up, the disaster site nodes are moved to the “Waiting for switchback” state.

If the disaster sites remained up, they are immediately moved to the “Waiting for switchback” state.

- The healing operations are performed automatically.

You can confirm the status of the disaster site nodes, and that healing operations succeeded, by using the `metrocluster operation show` command.



If automatic healing fails

If the automatic healing operation fails for any reason, you must issue the `metrocluster heal` commands manually as done in ONTAP versions prior to ONTAP 9.6. You can use the `metrocluster operation show` and `metrocluster operation history show -instance` commands to monitor the status of healing and determine the cause of a failure.

Creating SVMs for a MetroCluster configuration

You can create SVMs for a MetroCluster configuration to provide synchronous disaster recovery and high availability of data on clusters that are set up for a MetroCluster configuration.

- The two clusters must be in a MetroCluster configuration.
- Aggregates must be available and online in both clusters.
- If required, IPspaces with the same names must be created on both clusters.
- If one of the clusters forming the MetroCluster configuration is rebooted without utilizing a switchover, then the sync-source SVMs might come online as “stopped” instead of “started”.

When you create an SVM on one of the clusters in a MetroCluster configuration, the SVM is created as the source SVM, and the partner SVM is automatically created with the same name but with the “-mc” suffix on the partner cluster. If the SVM name contains a period, the “-mc” suffix is applied prior to the first period, for example, SVM-MC.DNS.NAME.

In a MetroCluster configuration, you can create 64 SVMs on a cluster. A MetroCluster configuration supports 128 SVMs.

1. Use the `vserver create` command.

The following example shows the SVM with the subtype “sync-source” on the local site and the SVM with the subtype “sync-destination” on the partner site:

```
cluster_A::>vserver create -vserver vs4 -rootvolume vs4_root -aggregate
aggr1
-rootvolume-security-style mixed
[Job 196] Job succeeded:
Vserver creation completed
```

The SVM “vs4” is created on the local site and the SVM “vs4-mc” is created on the partner site.

2. View the newly created SVMs.

- On the local cluster, verify the configuration state of SVMs:

```
metrocluster vserver show
```

The following example shows the partner SVMs and their configuration state:

```
cluster_A::> metrocluster vserver show
```

Cluster	Vserver	Partner Vserver	Configuration State
cluster_A	vs4	vs4-mc	healthy
cluster_B	vs1	vs1-mc	healthy

- From the local and partner clusters, verify the state of the newly configured SVMs:

`vserver show` command

The following example displays the administrative and operational states of the SVMs:

```
cluster_A::> vserver show

Vserver Type      Subtype           Admin   Operational  Root
-----  -----  -----  -----  -----  -----
vs4      data    sync-source   running   running   vs4_root   aggr1

cluster_B::> vserver show

Vserver Type      Subtype           Admin   Operational  Root
-----  -----  -----  -----  -----  -----
vs4-mc   data    sync-destination  running  stopped    vs4_root   aggr1
```

SVM creation might fail if any intermediate operations, such as root volume creation, fail and the SVM is in the “initializing” state. You must delete the SVM and re-create it.

The SVMs for the MetroCluster configuration are created with a root volume size of 1 GB. The sync-source SVM is in the “running” state, and the sync-destination SVM is in the “stopped” state.

What happens during a switchback

After the disaster site has recovered and aggregates have healed, the MetroCluster switchback process returns storage and client access from the disaster recovery site to the home cluster.

The `metrocluster switchback` command returns the primary site to full, normal MetroCluster operation. Any configuration changes are propagated to the original SVMs. Data server operation is then returned to the sync-source SVMs on the disaster site and the sync-dest SVMs that had been operating on the surviving site are deactivated.

If SVMs were deleted on the surviving site while the MetroCluster configuration was in switchover state, the switchback process does the following:

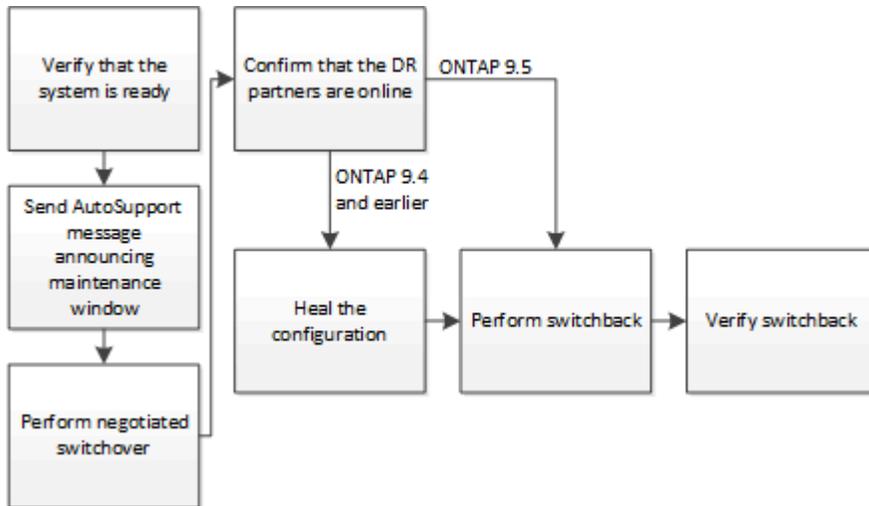
- Deletes the corresponding SVMs on the partner site (the former disaster site).
- Deletes any peering relationships of the deleted SVMs.

Perform switchover, healing, and switchback

Perform switchover for tests or maintenance

Performing switchover for tests or maintenance

If you want to test the MetroCluster functionality or to perform planned maintenance, you can perform a negotiated switchover in which one cluster is cleanly switched over to the partner cluster. You can then heal and switch back the configuration.



Beginning with ONTAP 9.6, switchover and switchback operations can be performed on MetroCluster IP configurations with ONTAP System Manager.

Limitations when the MetroCluster configuration is in switchover

When the system is in switchover, certain operations should not be performed. Learn about restricted operations when the system is in switchover.

Restricted operations in switchover

The following operations are not supported when the system is in switchover:

- Creating or deleting aggregates and volumes
- Creating or deleting SVMs
- Creating or deleting LIFs
- Adding or removing disks (only if you are replacing them as part of a recovery procedure)
- Performing configuration changes to SnapMirror SVM DR
- Modifying existing broadcast domains or creating new broadcast domains
- Modifying network subnets

Hardware replacement in switchover

Use the following procedures to replace controller hardware when the system is in switchover:

- If you need to replace a controller of the same type, at the site not in switchover, follow the procedure to [Recover from a multi-controller or storage failure](#).
 - If you need to replace the controller modules and chassis while the nodes are switched over at the surviving site, shut down both controllers and then perform the procedure to [Recover from a multi-controller or storage failure](#).
- If you need to replace a controller with a different type of controller, follow the procedure for your configuration in [Choose a controller upgrade procedure](#).
 - If your system is in switchover because of a controller failure or if you experience a controller failure while in switchover, you must first replace the controller hardware, perform a switchback, and then perform a controller upgrade:
 1. To replace the controller hardware and perform the switchback, follow [Recover from a multi-controller or storage failure](#).
 2. After you have replaced the hardware, perform a controller upgrade using the procedures in [Choose a controller upgrade procedure](#).

Verifying that your system is ready for a switchover

You can use the `-simulate` option to preview the results of a switchover operation. A verification check gives you a way to verify that most of the preconditions for a successful run are met before you start the operation. Issue these commands from the site that will remain up and operational:

1. Set the privilege level to advanced: `set -privilege advanced`
2. From the site that will remain up and operational, simulate a switchover operation: `metrocluster switchover -simulate`
3. Review the output that is returned.

The output shows whether any vetoes would prevent a switchover operation. Every time you perform a MetroCluster operation, you must verify a set of criteria for the success of the operation. A “veto” is a mechanism to prohibit the operation if one or more of the criteria are not fulfilled. There are two types of veto: a “soft” veto and a “hard” veto. You can override a soft veto, but not a hard veto. For example, to perform a negotiated switchover in a four-node MetroCluster configuration, one criterion is that all of the nodes are up and healthy. Suppose one node is down and was taken over by its HA partner. The switchover operation will be hard vetoed because it is a hard criterion that all of the nodes must be up and healthy. Because this is a hard veto, you cannot override the veto.



It is best not to override any veto.

Example: Verification results

The following example shows the errors that are encountered in a simulation of a switchover operation:

```
cluster4::*> metrocluster switchover -simulate
```

```
[Job 126] Preparing the cluster for the switchover operation...  
[Job 126] Job failed: Failed to prepare the cluster for the switchover  
operation. Use the "metrocluster operation show" command to view detailed  
error  
information. Resolve the errors, then try the command again.
```



Negotiated switchover and switchback will fail until you replace all of the failed disks. You can perform disaster recovery after you replace the failed disks. If you want to ignore the warning for failed disks, you can add a soft veto for the negotiated switchover and switchback.

Sending a custom AutoSupport message prior to negotiated switchover

Before performing a negotiated switchover, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. The negotiated switchover might result in plex or MetroCluster operation failures that trigger AutoSupport messages. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

This task must be performed on each MetroCluster site.

Steps

1. Log in to the cluster at Site_A.
2. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport  
invoke -node * -type all -message MAINT=maintenance-window-in-hours
```

`maintenance-window-in-hours` specifies the length of the maintenance window and can be a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can issue a command to indicating that the maintenance period has ended:

```
system node autosupport invoke -node *  
-type all -message MAINT=end
```
3. Repeat this step on the partner site.

Performing a negotiated switchover

A negotiated switchover cleanly shuts down processes on the partner site, and then switches over operations from the partner site. You can use a negotiated switchover to perform maintenance on a MetroCluster site or to test the switchover functionality.

- All previous configuration changes must be completed before performing a switchback operation.

This is to avoid competition with the negotiated switchover or switchback operation.

- Any nodes that were previously down must be booted and in cluster quorum.

The *System Administration Reference* has more information about cluster quorum in the “Understanding quorum and epsilon” section.

System administration

- The cluster peering network must be available from both sites.
- All of the nodes in the MetroCluster configuration must be running the same version of ONTAP software.
- The option replication.create_data_protection_rels.enable must be set to ON on both of the sites in a MetroCluster configuration before creating a new SnapMirror relationship.
- For a two-node MetroCluster configuration, a new SnapMirror relationship should not be created during an upgrade when there are mismatched versions of ONTAP between the sites.
- For a four-node MetroCluster configuration, the mismatched versions of ONTAP between the sites are not supported.

The recovering site can take a few hours to be able to perform the switchback operation.

The metrocluster switchover command switches over the nodes in all DR groups in the MetroCluster configuration. For example, in an eight-node MetroCluster configuration, it switches over the nodes in both DR groups.

While preparing for and executing a negotiated switchover, you must not make configuration changes to either cluster or perform any takeover or giveback operations.

For MetroCluster FC configurations:

- Mirrored aggregates will remain in normal state if the remote storage is accessible.
- Mirrored aggregates will become degraded after the negotiated switchover if access to the remote storage is lost.
- Unmirrored aggregates that are located at the disaster site will become unavailable if access to the remote storage is lost. This might lead to a controller outage.

For MetroCluster IP configurations:



Before performing maintenance tasks, you must remove monitoring if the MetroCluster configuration is monitored with the Tiebreaker or Mediator utility.

[Remove ONTAP Mediator or Tiebreaker monitoring before performing maintenance tasks](#)

- For ONTAP 9.4 and earlier:
 - Mirrored aggregates will become degraded after the negotiated switchover.
- For ONTAP 9.5 and later:
 - Mirrored aggregates will remain in normal state if the remote storage is accessible.
 - Mirrored aggregates will become degraded after the negotiated switchover if access to the remote storage is lost.
- For ONTAP 9.8 and later:
 - Unmirrored aggregates that are located at the disaster site will become unavailable if access to the remote storage is lost. This might lead to a controller outage.
 1. Use the metrocluster check run, metrocluster check show and metrocluster check config-replication show commands to make sure no configuration updates are in progress or pending. Issue these commands from the site that will remain up and operational.
 2. From the site that will remain up and operational, implement the switchover: `metrocluster switchover`

The operation can take several minutes to complete.

3. Monitor the completion of the switchover: `metrocluster operation show`

```
cluster_A::*> metrocluster operation show
  Operation: Switchover
  Start time: 10/4/2012 19:04:13
  State: in-progress
  End time: -
  Errors:

cluster_A::*> metrocluster operation show
  Operation: Switchover
  Start time: 10/4/2012 19:04:13
  State: successful
  End time: 10/4/2012 19:04:22
  Errors: -
```

4. Reestablish any SnapMirror or SnapVault configurations.

Verify that the SVMs are running and the aggregates are online

After the switchover is complete, you should verify that the DR partners have taken ownership of the disks and the partner SVMs have come online.

When you run the storage aggregate `plex show` command after a MetroCluster switchover, the status of `plex0` of the switched over root aggregate is indeterminate and is displayed as failed. During this time, the switched over root is not updated. The actual status of this plex can only be determined after the MetroCluster healing phase.

Steps

1. Check that the aggregates were switched over by using the `storage aggregate show` command.

In this example, the aggregates were switched over. The root aggregate (`aggr0_b2`) is in a degraded state. The data aggregate (`b2_aggr2`) is in a mirrored, normal state:

```

cluster_A::*> storage aggregate show

.
.
.
mcc1-b Switched Over Aggregates:
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
aggr0_b2      227.1GB   45.1GB   80% online    0 node_A_1
raid_dp,

mirror

degraded
b2_aggr1      227.1GB   200.3GB  20% online    0 node_A_1
raid_dp,

mirrored

normal

```

2. Confirm that the secondary SVMs have come online by using the vserver show command.

In this example, the previously dormant sync-destination SVMs on the secondary site have been activated and have an Admin State of running:

```

cluster_A::*> vserver show

Name      Name                               Admin      Operational  Root
Vserver   Type  Subtype                               State      State        Volume
Aggregate Service Mapping
-----
-----
...
cluster_B-vs1b-mc data  sync-destination  running    running
vs1b_vol  aggr_b1  file  file

```

Heal the configuration

Heal the configuration in a MetroCluster FC configuration

Healing the configuration in a MetroCluster FC configuration

Following a switchover, you must perform the healing operations in specific order to restore MetroCluster functionality.

- Switchover must have been performed and the surviving site must be serving data.
- Nodes on the disaster site must be halted or remain powered off.

They must not be fully booted during the healing process.

- Storage at the disaster site must be accessible (shelves are powered up, functional, and accessible).
- In fabric-attached MetroCluster configurations, inter-switch links (ISLs) must be up and operating.
- In four-node MetroCluster configurations, nodes in the surviving site must not be in HA failover state (all nodes must be up and running for each HA pair).

The healing operation must first be performed on the data aggregates, and then on the root aggregates.

Healing the data aggregates after negotiated switchover

You must heal the data aggregates after completing any maintenance or testing. This process resynchronizes the data aggregates and prepares the disaster site for normal operation. You must heal the data aggregates prior to healing the root aggregates.

All configuration updates in the remote cluster successfully replicate to the local cluster. You power up the storage on the disaster site as part of this procedure, but you do not and must not power up the controller modules on the disaster site.

Steps

1. Ensure that switchover has been completed by running the metrocluster operation show command.

```
controller_A_1::> metrocluster operation show
  Operation: switchover
    State: successful
  Start Time: 7/25/2014 20:01:48
  End Time: 7/25/2014 20:02:14
  Errors: -
```

2. Resynchronize the data aggregates by running the metrocluster heal -phase aggregates command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the --override -vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

3. Verify that the operation has been completed by running the metrocluster operation show command.

```

controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
  State: successful
Start Time: 7/25/2014 18:45:55
  End Time: 7/25/2014 18:45:56
  Errors: -

```

4. Check the state of the aggregates by running the storage aggregate show command.

```

controller_A_1::> storage aggregate show
Aggregate      Size Available Used% State   #Vols  Nodes           RAID
Status
-----
...
aggr_b2      227.1GB  227.1GB   0% online    0  mcc1-a2
raid_dp, mirrored, normal...

```

5. If storage has been replaced at the disaster site, you might need to remirror the aggregates.

Healing the root aggregates after negotiated switchover

After the data aggregates have been healed, you must heal the root aggregates in preparation for the switchback operation.

The data aggregates phase of the MetroCluster healing process must have been completed successfully.

Steps

1. Switch back the mirrored aggregates by running the metrocluster heal -phase root-aggregates command.

```

cluster_A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful

```

If the healing is vetoed, you have the option of reissuing the metrocluster heal command with the --override -vetoes parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

2. Confirm the heal operation is complete by running the metrocluster operation show command on the healthy cluster:

```
cluster_A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2014 20:54:41
  End Time: 7/29/2014 20:54:42
  Errors: -
```

3. Check for and remove any failed disks belonging to the disaster site by issuing the following command on the healthy site: `disk show -broken`
4. Power up or boot each controller module on the disaster site.

If the system displays the LOADER prompt, run the `boot_ontap` command.

5. After nodes are booted, verify that the root aggregates are mirrored.

If both plexes are present, resynchronization will occur automatically if the plexes are not synchronized. If one plex has failed, that plex must be destroyed and the mirror must be recreated using the `storage aggregate mirror -aggregateaggregate-name` command to reestablish the mirror relationship.

Healing the configuration in a MetroCluster IP configuration (ONTAP 9.4 and earlier)

You must heal the aggregates in preparation for the switchback operation.



On MetroCluster IP systems running ONTAP 9.5, healing is performed automatically, and you can skip these tasks.

The following conditions must exist before performing the healing procedure:

- Switchover must have been performed and the surviving site must be serving data.
- Storage shelves at the disaster site must be powered up, functional, and accessible.
- ISLs must be up and operating.
- Nodes in the surviving site must not be in HA failover state (both nodes must be up and running).

This task applies to MetroCluster IP configurations running ONTAP versions prior to 9.5 only.

This procedure differs from the healing procedure for MetroCluster FC configurations.

Steps

1. Power up each controller module on the site that was switched over and let them fully boot.

If the system displays the LOADER prompt, run the `boot_ontap` command.

2. Perform the root aggregate healing phase: `metrocluster heal root-aggregates`

```
cluster_A::> metrocluster heal root-aggregates
[Job 137] Job succeeded: Heal Root-Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal root-aggregates` command with the `--override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

3. Resynchronize the aggregates: `metrocluster heal aggregates`

```
cluster_A::> metrocluster heal aggregates
[Job 137] Job succeeded: Heal Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `--override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Confirm the heal operation is complete by running the `metrocluster operation show` command on the healthy cluster:

```
cluster_A::> metrocluster operation show
Operation: heal-aggregates
State: successful
Start Time: 7/29/2017 20:54:41
End Time: 7/29/2017 20:54:42
Errors: -
```

Performing a switchback

After you heal the MetroCluster configuration, you can perform the MetroCluster switchback operation. The MetroCluster switchback operation returns the configuration to its normal operating state, with the sync-source storage virtual machines (SVMs) on the disaster site active and serving data from the local disk pools.

- The disaster cluster must have successfully switched over to the surviving cluster.
- Healing must have been performed on the data and root aggregates.
- The surviving cluster nodes must not be in the HA failover state (all nodes must be up and running for each HA pair).
- The disaster site controller modules must be completely booted and not in the HA takeover mode.
- The root aggregate must be mirrored.
- The Inter-Switch Links (ISLs) must be online.
- Any required licenses must be installed on the system.

1. Confirm that all nodes are in the enabled state: `metrocluster node show`

The following example displays the nodes that are in the enabled state:

```

cluster_B::> metrocluster node show

DR                               Configuration  DR
Group Cluster Node              State          Mirroring Mode
-----
1      cluster_A
      node_A_1    configured    enabled    heal roots
completed
      node_A_2    configured    enabled    heal roots
completed
      cluster_B
      node_B_1    configured    enabled    waiting for
switchback recovery
      node_B_2    configured    enabled    waiting for
switchback recovery
4 entries were displayed.

```

2. Confirm that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations have been successfully completed: `metrocluster check lif show`
4. Perform a simulated switchback to verify the system is ready: `metrocluster switchback -simulate`
5. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```

cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results.
To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"

```

```
cluster_A::> metrocluster check show
Last Checked On: 9/13/2018 20:41:37
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok

6 entries were displayed.

6. Perform the switchback by running the metrocluster switchback command from any node in the surviving cluster: `metrocluster switchback`
7. Check the progress of the switchback operation: `metrocluster show`

The switchback operation is still in progress when the output displays waiting-for-switchback:

```
cluster_B::> metrocluster show
Cluster                Entry Name                State
-----
Local: cluster_B      Configuration state      configured
                      Mode                       switchover
                      AUSO Failure Domain     -
Remote: cluster_A    Configuration state      configured
                      Mode                       waiting-for-switchback
                      AUSO Failure Domain     -
```

The switchback operation is complete when the output displays normal:

```
cluster_B::> metrocluster show
Cluster                Entry Name                State
-----
Local: cluster_B      Configuration state      configured
                      Mode                       normal
                      AUSO Failure Domain     -
Remote: cluster_A    Configuration state      configured
                      Mode                       normal
                      AUSO Failure Domain     -
```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command. This command is at the advanced privilege level.

8. Reestablish any SnapMirror or SnapVault configurations.

In ONTAP 8.3, you need to manually reestablish a lost SnapMirror configuration after a MetroCluster switchback operation. In ONTAP 9.0 and later, the relationship is reestablished automatically.

Verifying a successful switchback

After performing the switchback, you want to confirm that all aggregates and storage virtual machines (SVMs) are switched back and online.

1. Verify that the switched-over data aggregates are switched back:

```
storage aggregate show
```

In the following example, aggr_b2 on node B2 has switched back:

```
node_B_1::> storage aggregate show
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
...
aggr_b2       227.1GB   227.1GB   0% online    0 node_B_2  raid_dp,
mirrored,
normal
```

2. Verify that all sync-destination SVMs on the surviving cluster are dormant (showing an operational state of “stopped”):

```
vserver show -subtype sync-destination
```

```
node_B_1::> vserver show -subtype sync-destination
Vserver      Type      Subtype      Admin      Operational  Root
Aggregate
-----
...
cluster_A-vs1a-mc data sync-destination
                                running      stopped     vs1a_vol   aggr_b2
```

Sync-destination aggregates in the MetroCluster configuration have the suffix “-mc” automatically appended to their name to help identify them.

3. Verify the sync-source SVMs on the disaster cluster are up and running:

```
vserver show -subtype sync-source
```

```
node_A_1::> vserver show -subtype sync-source
Vserver          Type      Subtype      Admin      Operational  Root
Aggregate
-----
...
vs1a             data      sync-source  running    running      vs1a_vol  aggr_b2
```

4. Confirm that the switchback operations succeeded by using the `metrocluster operation show` command.

If the command output shows...	Then...
That the switchback operation state is successful.	The switchback process is complete and you can proceed with operation of the system.
That the switchback operation or switchback-continuation-agent operation is partially successful.	Perform the suggested fix provided in the output of the <code>metrocluster operation show</code> command.

You must repeat the previous sections to perform the switchback in the opposite direction. If site_A did a switchover of site_B, have site_B do a switchover of site_A.

Commands for switchover, healing, and switchback

There are specific ONTAP commands for performing the MetroCluster disaster recovery processes.

If you want to...	Use this command...
Verify that switchover can be performed without errors or vetoes.	<code>metrocluster switchover -simulate</code> at the advanced privilege level
Verify that switchback can be performed without errors or vetoes.	<code>metrocluster switchback -simulate</code> at the advanced privilege level
Switch over to the partner nodes (negotiated switchover).	<code>metrocluster switchover</code>
Switch over to the partner nodes (forced switchover).	<code>metrocluster switchover -forced-on-disaster true</code>

Perform data aggregate healing.	<code>metrocluster heal -phase aggregates</code>
Perform root aggregate healing.	<code>metrocluster heal -phase root-aggregates</code>
Switch back to the home nodes.	<code>metrocluster switchback</code>

Use System Manger to perform switchover and switchback (MetroCluster IP configurations only)

You can switch over control from one MetroCluster IP site to the other to perform maintenance or recover from an issue.



Switchover and switchback procedures are supported only for MetroCluster IP configurations.

Overview of switchover and switchback

A switchover can occur in two instances:

- **A planned switchover**

This switchover is initiated by a system administrator using System Manager. The planned switchover allows a system administrator of a local cluster to switch control so that the data services of the remote cluster are handled by the local cluster. Then, a system administrator at the remote cluster location can perform maintenance on the remote cluster.

- **An unplanned switchover**

In some cases, when a MetroCluster cluster goes down or the connections between the clusters are down, ONTAP automatically initiates a switchover so that the cluster that is still running handles the data handling responsibilities of the cluster that is down.

At other times, when ONTAP cannot determine the status of one of the clusters, the system administrator of the site that is working initiates the switchover to take control of the data handling responsibilities of the other site.

For any type of switchover procedure, the data servicing capability is returned to the cluster by using a *switchback* process.

The switchover and switchback process you follow depends on your ONTAP version:

- [Use System Manager in ONTAP 9.6 or 9.7 for switchover and switchback](#)
- [Use System Manager in ONTAP 9.8 or later for switchover and switchback](#)

Use System Manager in ONTAP 9.6 or 9.7 for switchover and switchback

Steps

1. Log in to System Manager in ONTAP 9.6 or 9.7.
2. Click **(Return to classic version)**.

3. Click **Configuration > MetroCluster**.

System Manager verifies whether a negotiated switchover is possible.

4. Perform one of the following substeps when the validation process has completed:

- a. If validation fails, but Site B is up, then an error has occurred. For example, there might be a problem with a subsystem, or NVRAM mirroring might not be synchronized.
 - i. Fix the issue that is causing the error, click **Close**, and then start again at Step 2.
 - ii. Halt the Site B nodes, click **Close**, and then perform the steps in [Performing an unplanned switchover](#).
- b. If validation fails, and Site B is down, then there is most likely a connection problem. Verify that Site B is down, then perform the steps in [Performing an unplanned switchover](#).

5. Click **Switchover from Site B to Site A** to initiate the switchover process.

6. Click **Switch to the new experience**.

Use System Manager in ONTAP 9.8 or later for switchover and switchback

Perform a planned switchover (ONTAP 9.8 or later)

Steps

1. Log in to System Manager in ONTAP 9.8 or later.
2. Select **Dashboard**. In the **MetroCluster** section, the two clusters are shown with a connection.
3. In the local cluster (shown on the left), click , and select **Switchover remote data services to the local site**.

After the switchover request is validated, control is transferred from the remote site to the local site. The local site performs data service requests for both clusters.

The remote cluster reboots, but the storage components are not active, and the cluster does not service data requests. It is now available for planned maintenance.



The remote cluster should not be used for data servicing until you perform a switchback.

Perform an unplanned switchover (ONTAP 9.8 or later)

An unplanned switchover might be initiated automatically by ONTAP. If ONTAP cannot determine if a switchback is needed, the system administrator of the MetroCluster site that is still running initiates the switchover with the following steps:

Steps

1. Log in to System Manager in ONTAP 9.8 or later.
2. Select **Dashboard**.

In the **MetroCluster** section, the connection between the two clusters is shown with an "X" on it. This means that a connection cannot be detected and that either the connections or the cluster is down.

3. In the local cluster (shown on the left), click , and select **Switchover remote data services to the local site**.

If the switchover fails with an error, click on the "View details" link in the error message and confirm the unplanned switchover.

After the switchover request is validated, control is transferred from the remote site to the local site. The local site performs data service requests for both clusters.

The cluster must be repaired before it is brought online again.



After the remote cluster is brought online, it should not be used for data servicing until you perform a switchback.

Perform a switchback (ONTAP 9.8 or later)

Before you begin

If the remote cluster was down due to planned maintenance or due to a disaster, it should now be up and running and waiting for the switchback.

Steps

1. On the local cluster, log in to System Manager in ONTAP 9.8 or later.
2. Select **Dashboard**.

In the **MetroCluster** section, the two clusters are shown.

3. In the local cluster (shown on the left), click , and select **Take back control**.

The data is *healed* first, to verify the data is synchronized and mirrored between both clusters.

4. When the data healing is complete, click , and select **Initiate switchback**.

When the switchback is complete, both clusters are active and servicing data requests. Additionally, the data is being mirrored and synchronized between the clusters.

Monitoring the MetroCluster configuration

You can use ONTAP MetroCluster commands and Active IQ Unified Manager (formerly OnCommand Unified Manager) to monitor the health of a variety of software components and the state of MetroCluster operations.

Checking the MetroCluster configuration

You can check that the components and relationships in the MetroCluster configuration are working correctly. You should do a check after initial configuration and after making any changes to the MetroCluster configuration. You should also do a check before a negotiated (planned) switchover or a switchback operation.

About this task

If the `metrocluster check run` command is issued twice within a short time on either or both clusters, a conflict can occur and the command might not collect all data. Subsequent `metrocluster check show` commands do not show the expected output.

Steps

1. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```
cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results. To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"
```

2. Display more detailed results from the most recent metrocluster check run command:

```
metrocluster check aggregate show
```

```
metrocluster check cluster show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```

The metrocluster check show commands show the results of the most recent metrocluster check run command. You should always run the metrocluster check run command prior to using the metrocluster check show commands so that the information displayed is current.

The following example shows the metrocluster check aggregate show command output for a healthy four-node MetroCluster configuration:

```
cluster_A::> metrocluster check aggregate show

Last Checked On: 8/5/2014 00:42:58

Node          Aggregate          Check
Result
-----
controller_A_1  controller_A_1_aggr0
ok
ok
ok
ownership-state
```

```

ok
        controller_A_1_aggr1
                                mirroring-status
ok
                                disk-pool-allocation
ok
                                ownership-state
ok
        controller_A_1_aggr2
                                mirroring-status
ok
                                disk-pool-allocation
ok
                                ownership-state
ok

controller_A_2        controller_A_2_aggr0
                                mirroring-status
ok
                                disk-pool-allocation
ok
                                ownership-state
ok
        controller_A_2_aggr1
                                mirroring-status
ok
                                disk-pool-allocation
ok
                                ownership-state
ok
        controller_A_2_aggr2
                                mirroring-status
ok
                                disk-pool-allocation
ok
                                ownership-state
ok

18 entries were displayed.

```

The following example shows the `metrocluster check cluster show` command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

Last Checked On: 9/13/2017 20:47:04

Cluster	Check	Result
mccint-fas9000-0102	negotiated-switchover-ready	not-applicable
	switchback-ready	not-applicable
	job-schedules	ok
	licenses	ok
	periodic-check-enabled	ok
mccint-fas9000-0304	negotiated-switchover-ready	not-applicable
	switchback-ready	not-applicable
	job-schedules	ok
	licenses	ok
	periodic-check-enabled	ok

10 entries were displayed.

Commands for checking and monitoring the MetroCluster configuration

There are specific ONTAP commands for monitoring the MetroCluster configuration and checking MetroCluster operations.

Commands for checking MetroCluster operations

If you want to...	Use this command...
Perform a check of the MetroCluster operations. Note: This command should not be used as the only command for pre-DR operation system validation.	<code>metrocluster check run</code>
View the results of the last check on MetroCluster operations.	<code>metrocluster show</code>
View results of check on configuration replication between the sites.	<code>metrocluster check config-replication</code> <code>show metrocluster check config-replication</code> <code>show-aggregate-eligibility</code>
View results of check on node configuration.	<code>metrocluster check node show</code>
View results of check on aggregate configuration.	<code>metrocluster check aggregate show</code>
View the LIF placement failures in the MetroCluster configuration.	<code>metrocluster check lif show</code>

Commands for monitoring the MetroCluster interconnect

If you want to...	Use this command...
Display the HA and DR mirroring status and information for the MetroCluster nodes in the cluster.	<code>metrocluster interconnect mirror show</code>

Commands for monitoring MetroCluster SVMs

If you want to...	Use this command...
View all SVMs in both sites in the MetroCluster configuration.	<code>metrocluster vserver show</code>

Using the MetroCluster Tiebreaker or ONTAP Mediator to monitor the configuration

See [Differences between ONTAP Mediator and MetroCluster Tiebreaker](#) to understand the differences between these two methods of monitoring your MetroCluster configuration and initiating an automatic switchover.

Use these links to install and configure Tiebreaker or Mediator:

- [Install and configure the MetroCluster Tiebreaker software](#)
- [Prepare to install ONTAP Mediator](#)

How the NetApp MetroCluster Tiebreaker software detects failures

The Tiebreaker software resides on a Linux host. You need the Tiebreaker software only if you want to monitor two clusters and the connectivity status between them from a third site. Doing so enables each partner in a cluster to distinguish between an ISL failure, when inter-site links are down, from a site failure.

After you install the Tiebreaker software on a Linux host, you can configure the clusters in a MetroCluster configuration to monitor for disaster conditions.

How the Tiebreaker software detects intersite connectivity failures

The MetroCluster Tiebreaker software alerts you if all connectivity between the sites is lost.

Types of network paths

Depending on the configuration, there are three types of network paths between the two clusters in a MetroCluster configuration:

- **FC network (present in fabric-attached MetroCluster configurations)**

This type of network is composed of two redundant FC switch fabrics. Each switch fabric has two FC switches, with one switch of each switch fabric co-located with a cluster. Each cluster has two FC switches, one from each switch fabric. All of the nodes have FC (NV interconnect and FCP initiator) connectivity to each of the co-located IP switches. Data is replicated from cluster to cluster over the ISL.

- **Intercluster peering network**

This type of network is composed of a redundant IP network path between the two clusters. The cluster peering network provides the connectivity that is required to mirror the storage virtual machine (SVM) configuration. The configuration of all of the SVMs on one cluster is mirrored by the partner cluster.

- **IP network (present in MetroCluster IP configurations)**

This type of network is composed of two redundant IP switch networks. Each network has two IP switches, with one switch of each switch fabric co-located with a cluster. Each cluster has two IP switches, one from each switch fabric. All of the nodes have connectivity to each of the co-located FC switches. Data is replicated from cluster to cluster over the ISL.

Monitoring intersite connectivity

The Tiebreaker software regularly retrieves the status of intersite connectivity from the nodes. If NV interconnect connectivity is lost and the intercluster peering does not respond to pings, then the clusters assume that the sites are isolated and the Tiebreaker software triggers an alert as "AllLinksSevered". If a cluster identifies the "AllLinksSevered" status and the other cluster is not reachable through the network, then the Tiebreaker software triggers an alert as "disaster".

How the Tiebreaker software detects site failures

The NetApp MetroCluster Tiebreaker software checks the reachability of the nodes in a MetroCluster configuration and the cluster to determine whether a site failure has occurred. The Tiebreaker software also triggers an alert under certain conditions.

Components monitored by the Tiebreaker software

The Tiebreaker software monitors each controller in the MetroCluster configuration by establishing redundant connections through multiple paths to a node management LIF and to the cluster management LIF, both hosted on the IP network.

The Tiebreaker software monitors the following components in the MetroCluster configuration:

- Nodes through local node interfaces
- Cluster through the cluster-designated interfaces
- Surviving cluster to evaluate whether it has connectivity to the disaster site (NV interconnect, storage, and intercluster peering)

When there is a loss of connection between the Tiebreaker software and all of the nodes in the cluster and to the cluster itself, the cluster will be declared as "not reachable" by the Tiebreaker software. It takes around three to five seconds to detect a connection failure. If a cluster is unreachable from the Tiebreaker software, the surviving cluster (the cluster that is still reachable) must indicate that all of the links to the partner cluster are severed before the Tiebreaker software triggers an alert.



All of the links are severed if the surviving cluster can no longer communicate with the cluster at the disaster site through FC (NV interconnect and storage) and intercluster peering.

Failure scenarios during which Tiebreaker software triggers an alert

The Tiebreaker software triggers an alert when the cluster (all of the nodes) at the disaster site is down or unreachable and the cluster at the surviving site indicates the "AllLinksSevered" status.

The Tiebreaker software does not trigger an alert (or the alert is vetoed) in the following scenarios:

- In an eight-node MetroCluster configuration, if one HA pair at the disaster site is down
- In a cluster with all of the nodes at the disaster site down, one HA pair at the surviving site down, and the cluster at the surviving site indicates the "AllLinksSevered" status

The Tiebreaker software triggers an alert, but ONTAP vetoes that alert. In this situation, a manual switchover is also vetoed

- Any scenario in which the Tiebreaker software can either reach at least one node or the cluster interface at the disaster site, or the surviving site still can reach either node at the disaster site through either FC (NV interconnect and storage) or intercluster peering

How the ONTAP Mediator supports automatic unplanned switchover

[Learn about how the ONTAP Mediator supports automatic unplanned switchover in MetroCluster IP configurations.](#)

Monitoring and protecting the file system consistency using NVFAIL

The `-nvfail` parameter of the `volume modify` command enables ONTAP to detect nonvolatile RAM (NVRAM) inconsistencies when the system is booting or after a switchover operation. It also warns you and protects the system against data access and modification until the volume can be manually recovered.

If ONTAP detects any problems, database or file system instances stop responding or shut down. ONTAP then sends error messages to the console to alert you to check the state of the database or file system. You can enable NVFAIL to warn database administrators of NVRAM inconsistencies among clustered nodes that can compromise database validity.

After the NVRAM data loss during failover or boot recovery, NFS clients cannot access data from any of the nodes until the NVFAIL state is cleared. CIFS clients are unaffected.

How NVFAIL impacts access to NFS volumes or LUNs

The NVFAIL state is set when ONTAP detects NVRAM errors when booting, when a MetroCluster switchover operation occurs, or during an HA takeover operation if the NVFAIL option is set on the volume. If no errors are detected at startup, the file service is started normally. However, if NVRAM errors are detected or NVFAIL processing is enforced on a disaster switchover, ONTAP stops database instances from responding.

When you enable the NVFAIL option, one of the processes described in the following table takes place during bootup:

If...	Then...
ONTAP detects no NVRAM errors	File service starts normally.

<p>ONTAP detects NVRAM errors</p>	<ul style="list-style-type: none"> • ONTAP returns a stale file handle (ESTALE) error to NFS clients trying to access the database, causing the application to stop responding, crash, or shut down. <p>ONTAP then sends an error message to the system console and log file.</p> <ul style="list-style-type: none"> • When the application restarts, files are available to CIFS clients even if you have not verified that they are valid. <p>For NFS clients, files remain inaccessible until you reset the <code>in-nvfailed-state</code> option on the affected volume.</p>
<p>If one of the following parameters is used:</p> <ul style="list-style-type: none"> • <code>dr-force-nvfail</code> volume option is set • <code>force-nvfail-all</code> switchover command option is set. 	<p>You can unset the <code>dr-force-nvfail</code> option after the switchover, if the administrator is not expecting to force NVFAIL processing for possible future disaster switchover operations. For NFS clients, files remain inaccessible until you reset the <code>in-nvfailed-state</code> option on the affected volume.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> Using the <code>force-nvfail-all</code> option causes the <code>dr-force-nvfail</code> option to be set on all of the DR volumes processed during the disaster switchover.</p> </div>
<p>ONTAP detects NVRAM errors on a volume that contains LUNs</p>	<p>LUNs in that volume are brought offline. The <code>in-nvfailed-state</code> option on the volume must be cleared, and the NVFAIL attribute on the LUNs must be cleared by bringing each LUN in the affected volume online. You can perform the steps to check the integrity of the LUNs and recover the LUN from a Snapshot copy or backup as necessary. After all of the LUNs in the volume are recovered, the <code>in-nvfailed-state</code> option on the affected volume is cleared.</p>

Commands for monitoring data loss events

If you enable the NVFAIL option, you receive notification when a system crash caused by NVRAM inconsistencies or a MetroCluster switchover occurs.

By default, the NVFAIL parameter is not enabled.

If you want to...	Use this command...
-------------------	---------------------

Create a new volume with NVFAIL enabled	<code>volume create -nvfail on</code>
Enable NVFAIL on an existing volume	<code>volume modify</code> Note: You set the <code>-nvfail</code> option to "on" to enable NVFAIL on the created volume.
Display whether NVFAIL is currently enabled for a specified volume	<code>volume show</code> Note: You set the <code>-fields</code> parameter to "nvfail" to display the NVFAIL attribute for a specified volume.

Related information

See the man page for each command for more information.

Accessing volumes in NVFAIL state after a switchover

After a switchover, you must clear the NVFAIL state by resetting the `-in-nvfailed-state` parameter of the `volume modify` command to remove the restriction of clients to access data.

Before you begin

The database or file system must not be running or trying to access the affected volume.

About this task

Setting `-in-nvfailed-state` parameter requires advanced-level privilege.

Step

1. Recover the volume by using the `volume modify` command with the `-in-nvfailed-state` parameter set to `false`.

After you finish

For instructions about examining database file validity, see the documentation for your specific database software.

If your database uses LUNs, review the steps to make the LUNs accessible to the host after an NVRAM failure.

Related information

[Monitoring and protecting the files system consistency using NVFAIL](#)

Recovering LUNs in NVFAIL states after switchover

After a switchover, the host no longer has access to data on the LUNs that are in NVFAIL states. You must perform a number of actions before the database has access to the LUNs.

Before you begin

The database must not be running.

Steps

1. Clear the NVFAIL state on the affected volume that hosts the LUNs by resetting the `-in-nvfailed-state` parameter of the `volume modify` command.
2. Bring the affected LUNs online.
3. Examine the LUNs for any data inconsistencies and resolve them.

This might involve host-based recovery or recovery done on the storage controller using SnapRestore.

4. Bring the database application online after recovering the LUNs.

Where to find additional information

You can learn more about MetroCluster configuration and operation.

MetroCluster and miscellaneous information

Information	Subject
MetroCluster Documentation	<ul style="list-style-type: none"> • All MetroCluster information
NetApp Technical Report 4375: NetApp MetroCluster for ONTAP 9.3	<ul style="list-style-type: none"> • A technical overview of the MetroCluster configuration and operation. • Best practices for MetroCluster configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the configuration • Configuring the MetroCluster in ONTAP
MetroCluster Tiebreaker 1.21 software installation and configuration	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
Active IQ Unified Manager documentation NetApp Documentation: Product Guides and Resources	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance

Copy-based transition

- Transitioning data from 7-Mode storage systems to clustered storage systems

Maintain the MetroCluster components

Learn about MetroCluster maintenance

Learn how to prepare for MetroCluster maintenance tasks and choose the correct maintenance procedure for your configuration.

Prepare for maintenance tasks

Review the information in [Prepare for MetroCluster maintenance](#) before performing any maintenance procedures.



You must enable console logging and remove ONTAP Mediator or Tiebreaker monitoring before you perform maintenance tasks.

Maintenance procedures for different types of MetroCluster configurations

- If you have a MetroCluster IP configuration, review the procedures in [Maintenance procedures for MetroCluster IP configurations](#).
- If you have a MetroCluster FC configuration, review the procedures in [Maintenance procedures for MetroCluster FC configurations](#).
- If you cannot find the procedure in the specific section for your configuration, review the procedures in [Maintenance procedures for all MetroCluster configurations](#).

All other maintenance procedures

The following table provides links to procedures related to MetroCluster maintenance that are not located in the three sections listed above:

Component	MetroCluster type (FC or IP)	Task	Procedure
ONTAP software	Both	ONTAP software upgrade	Upgrade, revert, or downgrade

Controller module	Both	FRU replacement (including replacement controller modules, PCIe cards, FC-VI card, and so on)	ONTAP Hardware Systems Documentation
		 <p>Moving a storage controller module or NVRAM card among the MetroCluster storage systems is not supported.</p>	
		Upgrade and expansion	MetroCluster Upgrade and Expansion
		Transition from FC to IP connectivity	Transition from MetroCluster FC to MetroCluster IP
Drive shelf	FC	All other shelf maintenance procedures. The standard procedures can be used.	Maintain DS460C DS224C and DS212C disk shelves
	IP	<p>All shelf maintenance procedures. The standard procedures can be used.</p> <p>If adding shelves for an unmirrored aggregate, see Considerations when using unmirrored aggregates</p>	Maintain DS460C DS224C and DS212C disk shelves

Prepare for MetroCluster maintenance

Enable console logging before performing maintenance tasks

Enable console logging on your devices before performing maintenance tasks.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions before performing maintenance procedures:

- Leave AutoSupport enabled during maintenance.

- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Remove ONTAP Mediator or Tiebreaker monitoring before performing maintenance tasks

Before performing maintenance tasks, you must remove monitoring if the MetroCluster configuration is monitored with the Tiebreaker or Mediator utility.

Maintenance tasks include upgrading the controller platform, upgrading ONTAP, and performing a negotiated switchover and switchback.

Steps

1. Collect the output for the following command:

```
storage iscsi-initiator show
```

2. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

If you are using...	Use this procedure...
Tiebreaker	Removing MetroCluster Configurations in the <i>MetroCluster Tiebreaker Installation and Configuration content</i>
Mediator	Issue the following command from the ONTAP prompt: metrocluster configuration-settings mediator remove
Third-party applications	Refer to the product documentation.

3. After completing maintenance of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations in the <i>MetroCluster Tiebreaker Installation and Configuration</i> section.

If you are using...	Use this procedure
Mediator	Configure ONTAP Mediator from a MetroCluster IP configuration in the <i>MetroCluster IP Installation and Configuration</i> section.
Third-party applications	Refer to the product documentation.

MetroCluster failure and recovery scenarios

You should be aware of how the MetroCluster configuration responds to different failure events.



For additional information about recovery from node failures, see the section "Choosing the correct recovery procedure" in the [Recover from a disaster](#).

Event	Impact	Recovery
Single node failure	A failover is triggered.	The configuration recovers through a local takeover. RAID is not impacted. Review system messages and replace failed FRUs as necessary. ONTAP Hardware Systems Documentation
Two nodes fail at one site	Two nodes will fail only if automated switchover is enabled in the MetroCluster Tiebreaker software.	Manual unplanned switchover (USO) if automated switchover in MetroCluster Tiebreaker software is not enabled. ONTAP Hardware Systems Documentation
MetroCluster IP interface—failure of one port	The system is degraded. Additional port failure impacts HA mirroring.	The second port is used. Health Monitor generates an alert if the physical link to the port is broken. Review system messages and replace failed FRUs as necessary. ONTAP Hardware Systems Documentation

MetroCluster IP interface—failure of both ports	HA capability is impacted. RAID SyncMirror of the node stops syncing.	Immediate manual recovery is required as there is no HA takeover. Review system messages and replace failed FRUs as necessary. ONTAP Hardware Systems Documentation
Failure of one MetroCluster IP switch	No impact. Redundancy is provided through the second network.	Replace the failed switch as necessary. Replacing an IP switch
Failure of two MetroCluster IP switches that are in the same network	No impact. Redundancy is provided through the second network.	Replace the failed switch as necessary. Replacing an IP switch
Failure of two MetroCluster IP switches that are at one site	RAID SyncMirror of the node stops syncing. HA capability is impacted and the cluster goes out of quorum.	Replace the failed switch as necessary. Replacing an IP switch
Failure of two MetroCluster IP switches that are at different sites and not on the same network (diagonal failure)	RAID SyncMirror of the node stops syncing.	RAID SyncMirror of the node stops syncing. Cluster and HA capability are not impacted. Replace the failed switch as necessary. Replacing an IP switch

Using the Interoperability Matrix Tool to find MetroCluster information

When setting up the MetroCluster configuration, you can use the Interoperability Tool to ensure you are using supported software and hardware versions.

[NetApp Interoperability Matrix Tool](#)

After opening the Interoperability Matrix, you can use the Storage Solution field to select your MetroCluster solution.

You use the **Component Explorer** to select the components and ONTAP version to refine your search.

You can click **Show Results** to display the list of supported configurations that match the criteria.

Maintenance procedures for MetroCluster FC configurations

Modify a switch or ATTO bridge IP address for health monitoring

After modifying the IP addresses of MetroCluster FC back-end switches and ATTO bridges, you must replace the old health monitoring IP addresses with the new values.

- [Modify a switch IP address](#)
- [Modify an ATTO bridge IP address](#)

Modify a switch IP address

Replace the old health monitoring IP address of a MetroCluster FC back-end switch.

Before you begin

Refer to the switch vendor's documentation for your switch model to change the IP address on the switch before changing the health monitoring IP address.

Steps

1. Run the `::> storage switch show` command and in the output, note the switches that are reporting errors.
2. Remove the switch entries with old IP addresses:

```
::> storage switch remove -name switch_name
```

3. Add the switches with new IP addresses:

```
::> storage switch add -name switch_name -address new_IP_address -managed-by  
in-band
```

4. Verify the new IP addresses and confirm that there are no errors:

```
::> storage switch show
```

5. If required, refresh the entries:

```
::> set advanced
```

```
::*> storage switch refresh
```

```
::*> set admin
```

Modify an ATTO bridge IP address

Replace the old health monitoring IP address of an ATTO bridge.

Steps

1. Run the `::> storage bridge show` command and in the output, note the ATTO bridges that are reporting errors.
2. Remove the ATTO bridge entries with old IP addresses:

```
::> storage bridge remove -name ATTO_bridge_name
```

3. Add the ATTO bridges with new IP addresses:

```
::> storage bridge add -name ATTO_bridge_name -address new_IP_address -managed  
-by in-band
```

4. Verify the new IP addresses and confirm that there are no errors:

```
::> storage bridge show
```

5. If required, refresh the entries:

```
::> set advanced
```

```
::*> storage bridge refresh
```

```
::*> set admin
```

FC-to-SAS bridge maintenance

Support for FibreBridge 7600N bridges in MetroCluster configurations

The FibreBridge 7600N bridge is supported on ONTAP 9.5 and later as a replacement for the FibreBridge 7500N or 6500N bridge or when adding new storage to the MetroCluster configuration. The zoning requirements and restrictions regarding use of the bridge's FC ports are the same as that of the FibreBridge 7500N bridge.

[NetApp Interoperability Matrix Tool](#)



FibreBridge 6500N bridges are not supported in configurations running ONTAP 9.8 and later.

Use case	Zoning changes needed?	Restrictions	Procedure
Replacing a single FibreBridge 7500N bridge with a single FibreBridge 7600N bridge	No	The FibreBridge 7600N bridge must be configured exactly the same as the FibreBridge 7500N bridge.	Hot-swapping a FibreBridge 7500N with a 7600N bridge
Replacing a single FibreBridge 6500N bridge with a single FibreBridge 7600N bridge	No	The FibreBridge 7600N bridge must be configured exactly the same as the FibreBridge 6500N bridge.	Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge

Adding new storage through adding a new pair of FibreBridge 7600N bridges	Yes You must add storage zones for each of the FC ports of the new bridges.	You must have available ports on the FC switch fabric (in a fabric-attached MetroCluster configuration) or on the storage controllers (in a stretch MetroCluster configuration). Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.	Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system
---	--	---	---

Support for FibreBridge 7500N bridges in MetroCluster configurations

The FibreBridge 7500N bridge is supported as a replacement for the FibreBridge 6500N bridge or for when adding new storage to the MetroCluster configuration. The supported configurations have zoning requirements and restrictions regarding use of the bridge's FC ports and stack and storage shelf limits.



FibreBridge 6500N bridges are not supported in configurations running ONTAP 9.8 and later.

Use case	Zoning changes needed?	Restrictions	Procedure
Replacing a single FibreBridge 6500N bridge with a single FibreBridge 7500N bridge	No	The FibreBridge 7500N bridge must be configured exactly the same as the FibreBridge 6500N bridge, using a single FC port and attaching to a single stack. The second FC port on the FibreBridge 7500N must not be used.	Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge
Consolidating multiple stacks by replacing multiple pairs of FibreBridge 6500N bridges with a single pair of FibreBridge 7500N bridges	Yes	In this case, you take the FibreBridge 6500N bridges out of service and replace them with a single pair of FibreBridge 7500N bridges. Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks. At the end of the procedure, both the top and bottom of the stacks must be connected to corresponding ports on the FibreBridge 7500N bridges.	Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges

Use case	Zoning changes needed?	Restrictions	Procedure
Adding new storage through adding a new pair of FibreBridge 7500N bridges	Yes You must add storage zones for each of the FC ports of the new bridges.	You must have available ports on the FC switch fabric (in a fabric-attached MetroCluster configuration) or on the storage controllers (in a stretch MetroCluster configuration). Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.	Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system

Enabling IP port access on the FibreBridge 7600N bridge if necessary

If you are using an ONTAP version prior to 9.5, or otherwise plan to use out-of-band access to the FibreBridge 7600N bridge using telnet or other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV), you can enable the access services via the console port.

Unlike the ATTO FibreBridge 7500N bridge, the FibreBridge 7600N bridge is shipped with all IP port protocols and services disabled.

Beginning with ONTAP 9.5, *in-band management* of the bridges is supported. This means the bridges can be configured and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required and the bridge user interfaces are not required.

Beginning with ONTAP 9.8, *in-band management* of the bridges is supported by default and out-of-band SNMP management is deprecated.

This task is required if you are **not** using in-band management to manage the bridges. In this case, you need to configure the bridge via the Ethernet management port.

Steps

1. Access the bridge's console interface by connecting a serial cable to the serial port on the FibreBridge 7600N bridge.
2. Using the console, enable the access services, and then save the configuration:

```
set closeport none
```

```
saveconfiguration
```

The `set closeport none` command enables all access services on the bridge.

3. Disable a service, if desired, by issuing the `set closeport` and repeating the command as necessary until all desired services are disabled:

```
set closeport service
```

The `set closeport` command disables a single service at a time.

service can specify one of the following:

- expressnav
- ftp
- icmp
- quicknav
- snmp
- telnet

You can check whether a specific protocol is enabled or disabled by using the `get closeport` command.

4. If you are enabling SNMP, you must also issue the `set SNMP enabled` command:

```
set SNMP enabled
```

SNMP is the only protocol that requires a separate enable command.

5. Save the configuration:

```
saveconfiguration
```

Updating firmware on a FibreBridge bridge

The procedure for updating the bridge firmware depends on your bridge model and ONTAP version.

About this task

[Enable console logging](#) before performing this task.

Updating firmware on FibreBridge 7600N or 7500N bridges on configurations running ONTAP 9.4 and later

You might need to update the firmware on your FibreBridge bridges to ensure that you have the latest features or to resolve possible issues. This procedure should be used for FibreBridge 7600N or 7500N bridges on configurations running ONTAP 9.4 and later.

- The MetroCluster configuration must be operating normally.
- All of the FibreBridge bridges in the MetroCluster configuration must be up and operating.
- All of the storage paths must be available.
- You need the admin password and access to an HTTP, FTP, or Trivial File Transfer Protocol (TFTP) server.
- You must be using a supported firmware version.

[NetApp Interoperability Matrix Tool](#)

In the IMT, you can use the Storage Solution field to select your MetroCluster solution. You use the **Component Explorer** to select the components and ONTAP version to refine your search. You can click **Show Results** to display the list of supported configurations that match the criteria.

- You can use this task only on FibreBridge 7600N or 7500N bridges in configurations running ONTAP 9.4 or later.

- You must perform this task on each FibreBridge bridge in the MetroCluster configuration, so that all of the bridges are running the same firmware version.



This procedure is nondisruptive and takes approximately 30 minutes to complete.



Beginning with ONTAP 9.8, the `system bridge` command replaces the `storage bridge`. The following steps show the `system bridge` command, but if you're running a version earlier than ONTAP 9.8, you should use the `storage bridge` command.

Steps

1. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

“maintenance-window-in-hours” specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

2. Go to the ATTO FibreBridge page and select the appropriate firmware for your bridge.

[ATTO FibreBridge Firmware Download Page](#)

3. Review the Caution/MustRead and End User Agreement, and click the check box to indicate acceptance and proceed.
4. Place the firmware file in a network location that is network accessible to the controller modules.

You can enter the commands in the remaining steps from the console of either controller module.

5. Change to the advanced privilege level:

```
set -privilege advanced
```

You must respond with “y” when prompted to continue into advanced mode and see the advanced mode prompt (*>).

6. Update the bridge firmware.

Beginning in ONTAP 9.16.1, you can use credentials to update the bridge firmware if they are required by the server to download the firmware package.

If credentials are not required:

- a. Update the bridge firmware:

```
system bridge firmware update -bridge <name> -uri <URL-of-firmware-  
package>
```

Example

```
cluster_A> system bridge firmware update -bridge bridge_A_1a -uri  
http://192.168.132.97/firmware.ZBD
```

If credentials are required:

- a. Update the bridge firmware and specify the required user name:

```
system bridge firmware update -bridge <name> -uri <URL-of-  
firmware-package> -username <name>
```

- b. Enter the password when prompted in the output, as shown in the following example:

Example

```
cluster_A> system bridge firmware update -bridge bridge_A_1a -uri  
http://192.168.132.97/firmware.ZBD -username abc  
  
(system bridge)  
  
Enter the password:  
  
[Job 70] Job is queued: System bridge firmware update job.
```

7. Return to the admin privilege level:

```
set -privilege admin
```

8. Verify that the firmware upgrade is complete:

```
job show -name "<job_name>"
```

The following example shows that the job “system bridge firmware update” is still running:

```
cluster_A> job show -name "system bridge firmware update"
Owning
```

Job ID	Name	Vserver	Node	State
2246	job-name	cluster_A	node_A_1	Running

Description: System bridge firmware update job

After approximately 10 minutes, the new firmware is fully installed and the job state will be Success:

```
cluster_A> job show -name "system bridge firmware update"
```

Job ID	Name	Vserver	Node	State
2246	System bridge firmware update	cluster_A	node_A_1	Success

Description: System bridge firmware update job

9. Complete the steps according to whether in-band management is enabled and which version of ONTAP your system is running:

- If you are running ONTAP 9.4, in-band management is not supported and the command must be issued from the bridge console:
 - i. Run the `flashimages` command on the console of the bridge and confirm that the correct firmware versions are displayed.



The example shows that primary flash image shows the new firmware image, while the secondary flash image shows the old image.

```
flashimages

;Type Version
;=====
Primary 3.16 001H
Secondary 3.15 002S
Ready.
```

- i. Reboot the bridge by running the `firmwarerestart` command from the bridge.

- If you are running ONTAP 9.5 or later, in-band management is supported and the command can be issued from the cluster prompt:
- ii. Run the system bridge `run-cli -name <bridge_name> -command FlashImages` command.



The example shows that primary flash image shows the new firmware image, while the secondary flash image shows the old image.

```
cluster_A> system bridge run-cli -name ATTO_7500N_IB_1 -command
FlashImages

[Job 2257]

;Type          Version
;=====
Primary 3.16 001H
Secondary 3.15 002S
Ready.

[Job 2257] Job succeeded.
```

- iii. If necessary, restart the bridge:

```
system bridge run-cli -name ATTO_7500N_IB_1 -command FirmwareRestart
```



Beginning with ATTO firmware version 2.95 the bridge will restart automatically and this step is not required.

- 10. Verify that the bridge restarted correctly:

```
sysconfig
```

The system should be cabled for multipath high availability (both controllers have access through the bridges to the disk shelves in each stack).

```
cluster_A> node run -node cluster_A-01 -command sysconfig
NetApp Release 9.6P8: Sat May 23 16:20:55 EDT 2020
System ID: 1234567890 (cluster_A-01); partner ID: 0123456789 (cluster_A-
02)
System Serial Number: 200012345678 (cluster_A-01)
System Rev: A4
System Storage Configuration: Quad-Path HA
```

- 11. Verify that the FibreBridge firmware was updated:

```
system bridge show -fields fw-version,symbolic-name
```

```
cluster_A> system bridge show -fields fw-version,symbolic-name
name fw-version symbolic-name
-----
ATTO_20000010affeaffe 3.10 A06X bridge_A_1a
ATTO_20000010affeaffae 3.10 A06X bridge_A_1b
ATTO_20000010affeaffff 3.10 A06X bridge_A_2a
ATTO_20000010affeafffa 3.10 A06X bridge_A_2b
4 entries were displayed.
```

12. Verify the partitions are updated from the bridge's prompt:

```
flashimages
```

The primary flash image displays the new firmware image, while the secondary flash image displays the old image.

```
Ready.
flashimages

;Type          Version
;=====
  Primary      3.16 001H
  Secondary    3.15 002S

Ready.
```

13. Repeat steps 5 to 10 to ensure that both flash images are updated to the same version.

14. Verify that both flash images are updated to the same version.

```
flashimages
```

The output should show the same version for both partitions.

```
Ready.
flashimages

;Type          Version
;=====
  Primary      3.16 001H
  Secondary    3.16 001H

Ready.
```

15. Repeat steps 5 to 13 on the next bridge until all of the bridges in the MetroCluster configuration have been updated.

Replacing a single FC-to-SAS bridge

You can nondisruptively replace a bridge with a same model bridge or with a new model bridge.

Before you begin

You need the admin password and access to an FTP or SCP server.

About this task

This procedure is nondisruptive and takes approximately 60 minutes to complete.

This procedure uses the bridge CLI to configure and manage a bridge, and to update the bridge firmware and the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port. You can use other interfaces if they meet the requirements.

[Requirements for using other interfaces to configure and manage FibreBridge bridges](#)

Related information

[Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges](#)

Verifying storage connectivity

Before replacing bridges, you should verify bridge and storage connectivity. Familiarizing yourself with the command output enables you to subsequently confirm connectivity after making configuration changes.

About this task

You can issue these commands from the admin prompt of any of the controller modules in the MetroCluster configuration at the site undergoing maintenance.

Steps

1. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2017
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
```

```

slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60130
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        UTILITIES CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0Q9R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
brcd6505-fcs29:12.126L1527  : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
brcd6505-fcs29:12.126L1528  : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
brcd6505-fcs42:13.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
brcd6505-fcs42:6.126L0       : ATTO      FibreBridge6500N 1.61
FB6500N101167
brcd6505-fcs42:7.126L0       : ATTO      FibreBridge6500N 1.61
FB6500N102974
.
.
.
**<List of storage shelves visible to port\>**
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.
.
.

```

Hot-swapping a bridge with a replacement bridge of the same model

You can hot-swap a failed bridge with another bridge of the same model.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. If the old bridge is accessible, you can retrieve the configuration information.

If...	Then...
You are using IP management	Connect to the old bridge with a Telnet connection and copy the output of the bridge configuration.
You are using in-band management	Use the ONTAP CLI to retrieve the configuration information with the following commands: <pre>storage bridge run-cli -name <i>bridge-name</i> -command "info"</pre> <pre>storage bridge run-cli -name <i>bridge-name</i> -command "sasportlist"</pre>

- a. Enter the command:

```
storage bridge run-cli -name bridge_A1 -command "info"
```

```
info

Device Status           = Good
Unsaved Changes        = None
Device                  = "FibreBridge 7500N"
Serial Number           = FB7500N100000
Device Version          = 3.10
Board Revision          = 7
Build Number            = 007A
Build Type              = Release
Build Date              = "Aug 20 2019" 11:01:24
Flash Revision          = 0.02
Firmware Version        = 3.10
BCE Version (FPGA 1)    = 15
BAU Version (FPGA 2)    = 33
```

```
User-defined name      = "bridgeA1"
World Wide Name       = 20 00 00 10 86 A1 C7 00
MB of RAM Installed   = 512
FC1 Node Name         = 20 00 00 10 86 A1 C7 00
FC1 Port Name         = 21 00 00 10 86 A1 C7 00
FC1 Data Rate         = 16Gb
FC1 Connection Mode   = ptp
FC1 FW Revision       = 11.4.337.0
FC2 Node Name         = 20 00 00 10 86 A1 C7 00
FC2 Port Name         = 22 00 00 10 86 A1 C7 00
FC2 Data Rate         = 16Gb
FC2 Connection Mode   = ptp
FC2 FW Revision       = 11.4.337.0
SAS FW Revision       = 3.09.52
MP1 IP Address        = 10.10.10.10
MP1 IP Subnet Mask    = 255.255.255.0
MP1 IP Gateway        = 10.10.10.1
MP1 IP DHCP           = disabled
MP1 MAC Address       = 00-10-86-A1-C7-00
MP2 IP Address        = 0.0.0.0 (disabled)
MP2 IP Subnet Mask    = 0.0.0.0
MP2 IP Gateway        = 0.0.0.0
MP2 IP DHCP           = enabled
MP2 MAC Address       = 00-10-86-A1-C7-01
SNMP                  = enabled
SNMP Community String = public
PS A Status           = Up
PS B Status           = Up
Active Configuration  = NetApp
```

Ready.

b. Enter the command:

```
storage bridge run-cli -name bridge_A1 -command "sasportlist"
```

SASPortList

```
;Connector      PHY      Link      Speed      SAS Address
;=====
Device A        1        Up        6Gb        5001086000a1c700
Device A        2        Up        6Gb        5001086000a1c700
Device A        3        Up        6Gb        5001086000a1c700
Device A        4        Up        6Gb        5001086000a1c700
Device B        1        Disabled  12Gb       5001086000a1c704
Device B        2        Disabled  12Gb       5001086000a1c704
Device B        3        Disabled  12Gb       5001086000a1c704
Device B        4        Disabled  12Gb       5001086000a1c704
Device C        1        Disabled  12Gb       5001086000a1c708
Device C        2        Disabled  12Gb       5001086000a1c708
Device C        3        Disabled  12Gb       5001086000a1c708
Device C        4        Disabled  12Gb       5001086000a1c708
Device D        1        Disabled  12Gb       5001086000a1c70c
Device D        2        Disabled  12Gb       5001086000a1c70c
Device D        3        Disabled  12Gb       5001086000a1c70c
Device D        4        Disabled  12Gb       5001086000a1c70c
```

2. If the bridge is in a fabric-attached MetroCluster configuration, disable all of the switch ports that connect to the bridge FC port or ports.
3. From the ONTAP cluster prompt, remove the bridge undergoing maintenance from health monitoring:
 - a. Remove the bridge:

```
storage bridge remove -name bridge-name
```

- b. View the list of monitored bridges and confirm that the removed bridge is not present:

```
storage bridge show
```

4. Properly ground yourself.
5. Power down the ATTO bridge and remove the power cables connected to the bridge.
6. Disconnect the cables that are connected to the old bridge.

You should make note of the port to which each cable was connected.

7. Remove the old bridge from the rack.
8. Install the new bridge into the rack.
9. Reconnect the power cord and, if configuring for IP access to the bridge, a shielded Ethernet cable.



You must not reconnect the SAS or FC cables at this time.

10. Connect the bridge to a power source, and then turn it on.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed

its power-on self test sequence.

11. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

12. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

13. Configure the bridge.

If you retrieved the configuration information from the old bridge, use the information to configure the new bridge.

Be sure to make note of the user name and password that you designate.

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

To set the IP address without the QuickNAV utility, you need to have a serial connection to the FibreBridge.

If using the CLI, you must run the following commands:

```
set ipaddress mp1 _ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
```

```
set ethernetspeed mp1 1000
```

- b. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

14. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

b. If you are configuring a FibreBridge 7500N, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode port-number ptp
```

c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable port-number
```

- d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used. If only SAS port A is used, then SAS ports B, C, and D must be disabled.

15. Secure access to the bridge and save the bridge's configuration.

- a. From the controller prompt check the status of the bridges: `storage bridge show`

The output shows which bridge is not secured.

- b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

- c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

- d. Save the bridge's configuration.

You must run the following commands:

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

16. Connect the FC cables to the same ports on the new bridge.

17. Update the FibreBridge firmware on each bridge.

If the new bridge is the same type as the partner bridge, upgrade to the same firmware as the partner bridge. If the new bridge is a different type to the partner bridge, upgrade to the latest firmware supported by the bridge and version of ONTAP. See [Updating firmware on a FibreBridge bridge](#)

18. Reconnect the SAS cables to the same ports on the new bridge.

You must replace the cables connecting the bridge to the top or bottom of the shelf stack. The FibreBridge 7600N and 7500N bridges require mini-SAS cables for these connections.



Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector). For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

- Verify that each bridge can see all of the disk drives and disk shelves to which the bridge is connected.

If you are using the...	Then...
ATTO ExpressNAV GUI	<ol style="list-style-type: none"> In a supported web browser, enter the IP address of the bridge in the browser box. You are brought to the ATTO FibreBridge homepage, which has a link. Click the link, and then enter your user name and the password that you designated when you configured the bridge. The ATTO FibreBridge status page appears with a menu to the left. Click Advanced in the menu. View the connected devices: <code>sastargets</code> Click Submit.
Serial port connection	View the connected devices: <code>sastargets</code>

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.



If the text response truncated appears at the beginning of the output, you can use Telnet to connect to the bridge, and then view all of the output by using the `sastargets` command.

The following output shows that 10 disks are connected:

```

Tgt VendorID ProductID      Type SerialNumber
  0 NETAPP    X410_S15K6288A15 DISK 3QP1CLE300009940UHJV
  1 NETAPP    X410_S15K6288A15 DISK 3QP1ELF600009940V1BV
  2 NETAPP    X410_S15K6288A15 DISK 3QP1G3EW00009940U2M0
  3 NETAPP    X410_S15K6288A15 DISK 3QP1EWMP00009940U1X5
  4 NETAPP    X410_S15K6288A15 DISK 3QP1FZLE00009940G8YU
  5 NETAPP    X410_S15K6288A15 DISK 3QP1FZLF00009940TZKZ
  6 NETAPP    X410_S15K6288A15 DISK 3QP1CEB400009939MGXL
  7 NETAPP    X410_S15K6288A15 DISK 3QP1G7A900009939FNNT
  8 NETAPP    X410_S15K6288A15 DISK 3QP1FY0T00009940G8PA
  9 NETAPP    X410_S15K6288A15 DISK 3QP1FXW600009940VERQ

```

20. Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat Step 19 for each remaining bridge.
Not correct	<ul style="list-style-type: none"> a. Check for loose SAS cables or correct the SAS cabling by repeating Step 18. b. Repeat Step 19.

21. If the bridge is in a fabric-attached MetroCluster configuration, re-enable the FC switch port that you disabled at the beginning of this procedure.

This should be the port that connects to the bridge.

22. From the system console of both controller modules, verify that all of the controller modules have access through the new bridge to the disk shelves (that is, that the system is cabled for Multipath HA):

```
run local sysconfig
```



It might take up to a minute for the system to complete discovery.

If the output does not indicate Multipath HA, you must correct the SAS and FC cabling because not all of the disk drives are accessible through the new bridge.

The following output states that the system is cabled for Multipath HA:

```
NetApp Release 8.3.2: Tue Jan 26 01:41:49 PDT 2016
System ID: 1231231231 (node_A_1); partner ID: 4564564564 (node_A_2)
System Serial Number: 700000123123 (node_A_1); partner Serial Number:
700000456456 (node_A_2)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: NA
```



When the system is not cabled as Multipath HA, restarting a bridge might cause loss of access to the disk drives and result in a multi-disk panic.

23. If running ONTAP 9.4 or earlier, verify that the bridge is configured for SNMP.

If you are using the bridge CLI, run the following command:

```
get snmp
```

24. From the ONTAP cluster prompt, add the bridge to health monitoring:

- a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

- b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```

controller_A_1::> storage bridge show

Bridge                Symbolic Name Is Monitored  Monitor Status
Vendor Model          Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01          true          ok          Atto
FibreBridge 7500N    20000010867038c0
ATTO_10.10.20.11  atto02          true          ok          Atto
FibreBridge 7500N    20000010867033c0
ATTO_10.10.20.12  atto03          true          ok          Atto
FibreBridge 7500N    20000010867030c0
ATTO_10.10.20.13  atto04          true          ok          Atto
FibreBridge 7500N    2000001086703b80

4 entries were displayed

controller_A_1::>

```

25. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Related information

[In-band management of the FC-to-SAS bridges](#)

Hot-swapping a FibreBridge 7500N with a 7600N bridge

You can hot-swap a FibreBridge 7500N bridge with a 7600N bridge.

About this task

If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. If the bridge is in a fabric-attached MetroCluster configuration, disable all of the switch ports that connect to the bridge FC port or ports.
2. From the ONTAP cluster prompt, remove the bridge undergoing maintenance from health monitoring:

- a. Remove the bridge:

```
storage bridge remove -name bridge-name
```

- b. View the list of monitored bridges and confirm that the removed bridge is not present:

```
storage bridge show
```

3. Properly ground yourself.
4. Remove the power cables connected to the bridge to power down the bridge.
5. Disconnect the cables that are connected to the old bridge.

You should make note of the port to which each cable was connected.

6. Remove the old bridge from the rack.
7. Install the new bridge into the rack.
8. Reconnect the power cord and shielded Ethernet cable.



You must not reconnect the SAS or FC cables at this time.

9. Connect the bridge to a power source, and then turn it on.

The bridge Ready LED might take up to 30 seconds to illuminate, indicating that the bridge has completed its power-on self test sequence.

10. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

11. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial

(COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

12. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

13. Configure the bridges.

Be sure to make note of the user name and password that you designate.

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on FibreBridge 7600N. The time synchronization for FibreBridge 7600N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

To set the IP address without the QuickNAV utility, you need to have a serial connection to the FibreBridge.

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
```

```
set ethernetspeed mp1 1000
```

- b. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b

- `bridge_B_1a`
- `bridge_B_1b`

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

- c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

14. Configure the bridge FC ports.

- a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the FC port of the controller module or switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

- b. You must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the unused port:

```
FCPortDisable port-number
```

The following example shows the disabling of FC port 2:

```
FCPortDisable 2

Fibre Channel Port 2 has been disabled.
```

- c. Disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used.

If only SAS port A is used, then SAS ports B, C, and D must be disabled. The following example shows disabling of SAS port B. You must similarly disable SAS ports C and D:

```
SASPortDisable b

SAS Port B has been disabled.
```

15. Secure access to the bridge and save the bridge's configuration.

- a. From the controller prompt check the status of the bridges:

```
storage bridge show
```

The output shows which bridge is not secured.

- b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

- c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

- d. Save the bridge's configuration.

You must run the following commands: +

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

16. Connect the FC cables to the same ports on the new bridge.

17. Update the FibreBridge firmware on each bridge.

[Update firmware on a FibreBridge bridge](#)

18. Reconnect the SAS cables to the same ports on the new bridge.



Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector). For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

- Verify that each bridge can see all of the disk drives and disk shelves to which the bridge is connected:

```
sastargets
```

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.

The following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNTT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

- Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat the previous step for each remaining bridge.
Not correct	<ol style="list-style-type: none"> Check for loose SAS cables or correct the SAS cabling by repeating Step 18. Repeat the previous step.

- If the bridge is in a fabric-attached MetroCluster configuration, reenab the FC switch port that you disabled at the beginning of this procedure.

This should be the port that connects to the bridge.

- From the system console of both controller modules, verify that all of the controller modules have access through the new bridge to the disk shelves (that is, that the system is cabled for Multipath HA):

```
run local sysconfig
```



It might take up to a minute for the system to complete discovery.

If the output does not indicate Multipath HA, you must correct the SAS and FC cabling because not all of the disk drives are accessible through the new bridge.

The following output states that the system is cabled for Multipath HA:

```
NetApp Release 8.3.2: Tue Jan 26 01:41:49 PDT 2016
System ID: 1231231231 (node_A_1); partner ID: 4564564564 (node_A_2)
System Serial Number: 700000123123 (node_A_1); partner Serial Number:
700000456456 (node_A_2)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: NA
```



When the system is not cabled as Multipath HA, restarting a bridge might cause loss of access to the disk drives and result in a multi-disk panic.

23. If running ONTAP 9.4 or earlier, verify that the bridge is configured for SNMP.

If you are using the bridge CLI, run the following command:

```
get snmp
```

24. From the ONTAP cluster prompt, add the bridge to health monitoring:

a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```

controller_A_1::> storage bridge show

Bridge                Symbolic Name Is Monitored  Monitor Status
Vendor Model          Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01          true           ok           Atto
FibreBridge 7500N    20000010867038c0
ATTO_10.10.20.11  atto02          true           ok           Atto
FibreBridge 7500N    20000010867033c0
ATTO_10.10.20.12  atto03          true           ok           Atto
FibreBridge 7500N    20000010867030c0
ATTO_10.10.20.13  atto04          true           ok           Atto
FibreBridge 7500N    2000001086703b80

4 entries were displayed

controller_A_1::>

```

25. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check: +

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Related information

In-band management of the FC-to-SAS bridges

Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge

You can hot-swap a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge to replace a failed bridge or upgrade your bridge in a fabric-attached or a bridge-attached MetroCluster configuration.

About this task

- This procedure is for hot-swapping a single FibreBridge 6500N bridge with single FibreBridge 7600N or 7500N bridge.
- When you hot-swap a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge, you must use only one FC port and one SAS port on the FibreBridge 7600N or 7500N bridge.
- If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



If you are hot-swapping both FibreBridge 6500N bridges in a pair, you must use the [Consolidate Multiple Storage Stacks](#) procedure for zoning instructions. By replacing both FibreBridge 6500N bridges on the bridge, you can take advantage of the additional ports on the FibreBridge 7600N or 7500N bridge.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. Do one of the following:
 - If the failed bridge is in a fabric-attached MetroCluster configuration, disable the switch port that connects to the bridge FC port.
 - If the failed bridge is in a stretch MetroCluster configuration, use either one of the available FC ports.
2. From the ONTAP cluster prompt, remove the bridge undergoing maintenance from health monitoring:
 - a. Remove the bridge:

```
storage bridge remove -name bridge-name
```

- b. View the list of monitored bridges and confirm that the removed bridge is not present:

```
storage bridge show
```

3. Properly ground yourself.
4. Turn off the power switch of the bridge.
5. Disconnect the cables connected from the shelf to the FibreBridge 6500N bridge ports and power cables.

You should make note of the ports that each cable was connected to.

6. Remove the FibreBridge 6500N bridge that you need to replace from the rack.
7. Install the new FibreBridge 7600N or 7500N bridge into the rack.
8. Reconnect the power cord and, if necessary, the shielded Ethernet cable.



Do not reconnect the SAS or FC cables at this time.

9. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

10. If configuring for IP management, connect the Ethernet management 1 port on each bridge to your network by using an Ethernet cable.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

The Ethernet management 1 port enables you to quickly download the bridge firmware (using ATTO ExpressNAV or FTP management interfaces) and to retrieve core files and extract logs.

11. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

12. Configure the bridge.

If you retrieved the configuration information from the old bridge, use the information to configure the new bridge.

Be sure to make note of the user name and password that you designate.

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

To set the IP address without the QuickNAV utility, you need to have a serial connection to the FibreBridge.

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address
```

```
set ipsubnetmask mp1 subnet-mask

set ipgateway mp1 x.x.x.x

set ipdhcp mp1 disabled

set ethernetspeed mp1 1000
```

b. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

13. Configure the bridge FC ports.

a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.
- The FibreBridge 6500N bridge supports up to 8, 4, or 2 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

b. If you are configuring a FibreBridge 7500N or 6500N bridge, configure the connection mode that the port uses to ptp.



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode port-number ptp
```

c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable port-number
```

d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used. If only SAS port A is used, then SAS ports B, C, and D must be disabled.

14. Secure access to the bridge and save the bridge's configuration.

a. From the controller prompt check the status of the bridges:

```
storage bridge show
```

The output shows which bridge is not secured.

b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

d. Save the bridge's configuration.

You must run the following commands:

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

15. Turn on Health Monitoring for the FibreBridge 7600N or 7500N bridge.

16. Connect the FC cables to the Fibre Channel 1 ports on the new bridge.

You must cable the FC port to the same switch or controller port that the FibreBridge 6500N bridge had been connected to.

17. Update the FibreBridge firmware on each bridge.

If the new bridge is the same type as the partner bridge, upgrade to the same firmware as the partner bridge. If the new bridge is a different type to the partner bridge, upgrade to the latest firmware and version of ONTAP supported by the bridge.

[Update firmware on a FibreBridge bridge](#)

18. Reconnect the SAS cables to the SAS A ports on the new bridge.

The SAS port must be cabled to the same shelf port that the FibreBridge 6500N bridge had been connected to.



Do not force a connector into a port. The mini-SAS cables are keyed; when oriented correctly into a SAS port, the SAS cable clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector). For controllers, the orientation of SAS ports can vary depending on the platform model; therefore, the correct orientation of the SAS cable connector varies.

19. Verify that the bridge can detect all of the disk drives and disk shelves it is connected to.

If you are using the...	Then...
ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of the bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge homepage, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page appears with a menu to the left.</p> <p>c. Click Advanced in the menu.</p> <p>d. Enter the following command and then click Submit to see the list of disks visible to the bridge:</p> <pre>sastargets</pre>
Serial port connection	<p>Display the list of disks visible to the bridge:</p> <pre>sastargets</pre>

The output shows the devices (disks and disk shelves) that the bridge is connected to. Output lines are sequentially numbered so that you can quickly count the devices. For example, the following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ



If the text “response truncated” appears at the beginning of the output, you can use Telnet to access the bridge and enter the same command to see all of the output.

20. Verify that the command output shows that the bridge is connected to all of the necessary disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat the previous step for each remaining bridge.
Not correct	<ul style="list-style-type: none"> a. Check for loose SAS cables or correct the SAS cabling by repeating Step 18. b. Repeat the previous step for each remaining bridge.

21. Reenable the FC switch port that connects to the bridge.
22. Verify that all controllers have access through the new bridge to the disk shelves (that the system is cabled for Multipath HA), at the system console of both controllers:

```
run local sysconfig
```



It might take up to a minute for the system to complete discovery.

For example, the following output shows that the system is cabled for Multipath HA:

```
NetApp Release 8.3.2: Tue Jan 26 01:23:24 PST 2016
System ID: 1231231231 (node_A_1); partner ID: 4564564564 (node_A_2)
System Serial Number: 700000123123 (node_A_1); partner Serial Number:
700000456456 (node_A_2)
System Rev: B0
System Storage Configuration: Multi-Path HA
System ACP Connectivity: NA
```

If the command output indicates that the configuration is mixed-path or single-path HA, you must correct

the SAS and FC cabling because not all disk drives are accessible through the new bridge.



When the system is not cabled as Multipath HA, restarting a bridge might cause loss of access to the disk drives and result in a multi-disk panic.

23. From the ONTAP cluster prompt, add the bridge to health monitoring:

a. Add the bridge, using the command for your version of ONTAP:

ONTAP version	Command
9.5 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>
9.4 and earlier	<code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>

b. Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the “Status” column is “ok”, and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```
controller_A_1::> storage bridge show

Bridge          Symbolic Name Is Monitored  Monitor Status
Vendor Model    Bridge WWN
-----
-----
ATTO_10.10.20.10  atto01      true         ok           Atto
FibreBridge 7500N  20000010867038c0
ATTO_10.10.20.11  atto02      true         ok           Atto
FibreBridge 7500N  20000010867033c0
ATTO_10.10.20.12  atto03      true         ok           Atto
FibreBridge 7500N  20000010867030c0
ATTO_10.10.20.13  atto04      true         ok           Atto
FibreBridge 7500N  2000001086703b80

4 entries were displayed

controller_A_1::>
```

24. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

25. After replacing the part, return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return & Replacements](#) page for further information.

Related information

[In-band management of the FC-to-SAS bridges](#)

Replacing a pair of FibreBridge 6500N bridges with 7600N or 7500N bridges

To take advantage of the additional FC2 port on the FibreBridge 7600N or 7500N bridges and reduce rack utilization, you can nondisruptively replace 6500N bridges and consolidate up to four storage stacks behind a single pair of FibreBridge 7600N or 7500N bridges.

Before you begin

You need the admin password and access to an FTP or SCP server.

About this task

You should use this procedure if:

- You are replacing a pair of FibreBridge 6500N bridges with FibreBridge 7600N or 7500N bridges.

After the replacement, both bridges in the pair must be the same model.

- You previously replaced a single FibreBridge 6500N bridge with a 7600N or 7500N bridge and are now replacing the second bridge in the pair.
- You have a pair of FibreBridge 7600N or 7500N bridges with available SAS ports and you are consolidating SAS storage stacks that are currently connected using FibreBridge 6500N bridges.

This procedure is nondisruptive and takes approximately two hours to complete.

Related information

[Replacing a single FC-to-SAS bridge](#)

Verifying storage connectivity

Before replacing bridges, you should verify bridge and storage connectivity. Familiarizing yourself with the command output enables you to subsequently confirm connectivity after making configuration changes.

You can issue these commands from the admin prompt of any of the controller modules in the MetroCluster configuration at the site undergoing maintenance.

1. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2017
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60130
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        UTILITIES CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0Q9R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
```

```

      ID      Vendor      Model      FW      Size
      brcd6505-fcs29:12.126L1527      : NETAPP      X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
      brcd6505-fcs29:12.126L1528      : NETAPP      X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
      .
      .
      .
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
      brcd6505-fcs40:12.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
      brcd6505-fcs42:13.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
      brcd6505-fcs42:6.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N101167
      brcd6505-fcs42:7.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102974
      .
      .
      .
**<List of storage shelves visible to port\>**
      brcd6505-fcs40:12.shelf6: DS4243      Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      brcd6505-fcs40:12.shelf8: DS4243      Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      .
      .
      .

```

Hot-swapping FibreBridge 6500N bridges to create a pair of FibreBridge 7600N or 7500N bridges

To hot-swap one or two FibreBridge 6500N bridges to create a configuration with a pair of FibreBridge 7600N or 7500N bridges, you must replace the bridges one at a time and follow the correct cabling procedure. The new cabling is different from the original cabling.

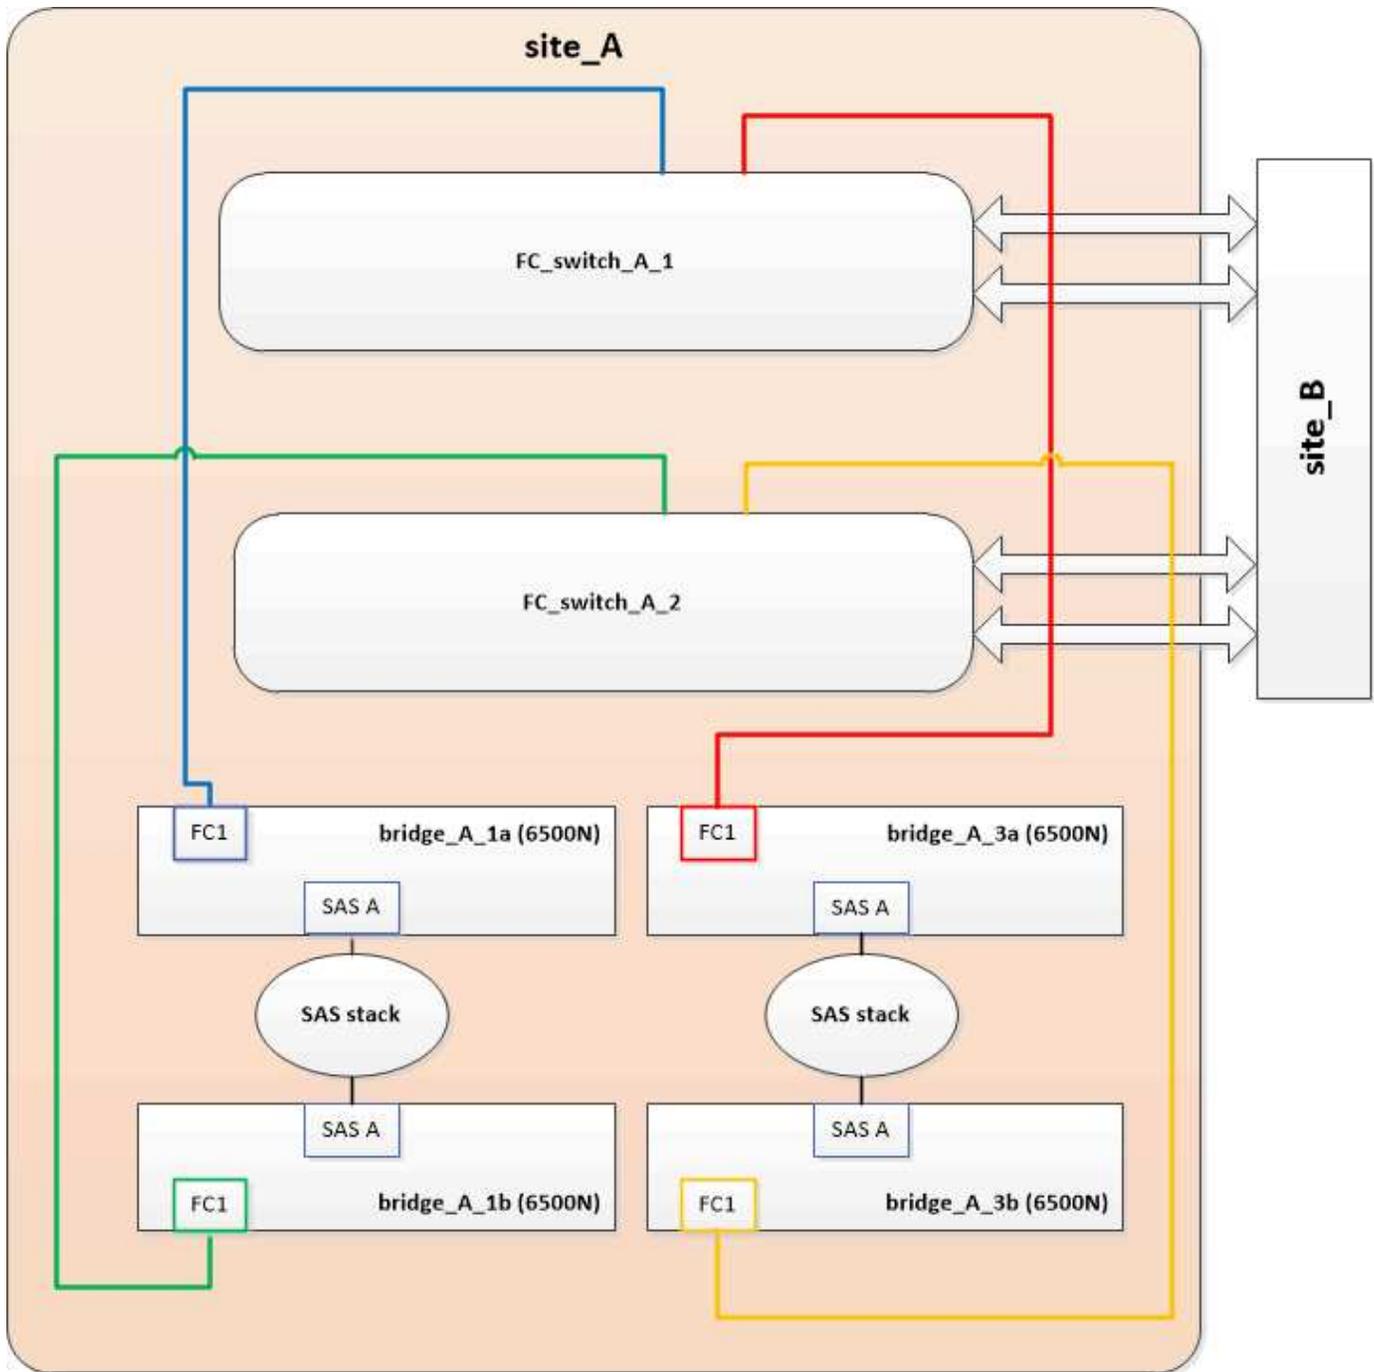
About this task

You can also use this procedure if the following conditions are true:

- You are replacing a pair of FibreBridge 6500N bridges that are both connected to the same stack of SAS storage.
- You previously replaced one FibreBridge 6500N bridge in the pair, and your storage stack is configured with one FibreBridge 6500N bridge and one FibreBridge 7600N or 7500N bridge.

In this case, you should start with the step below to hot-swap the bottom FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge.

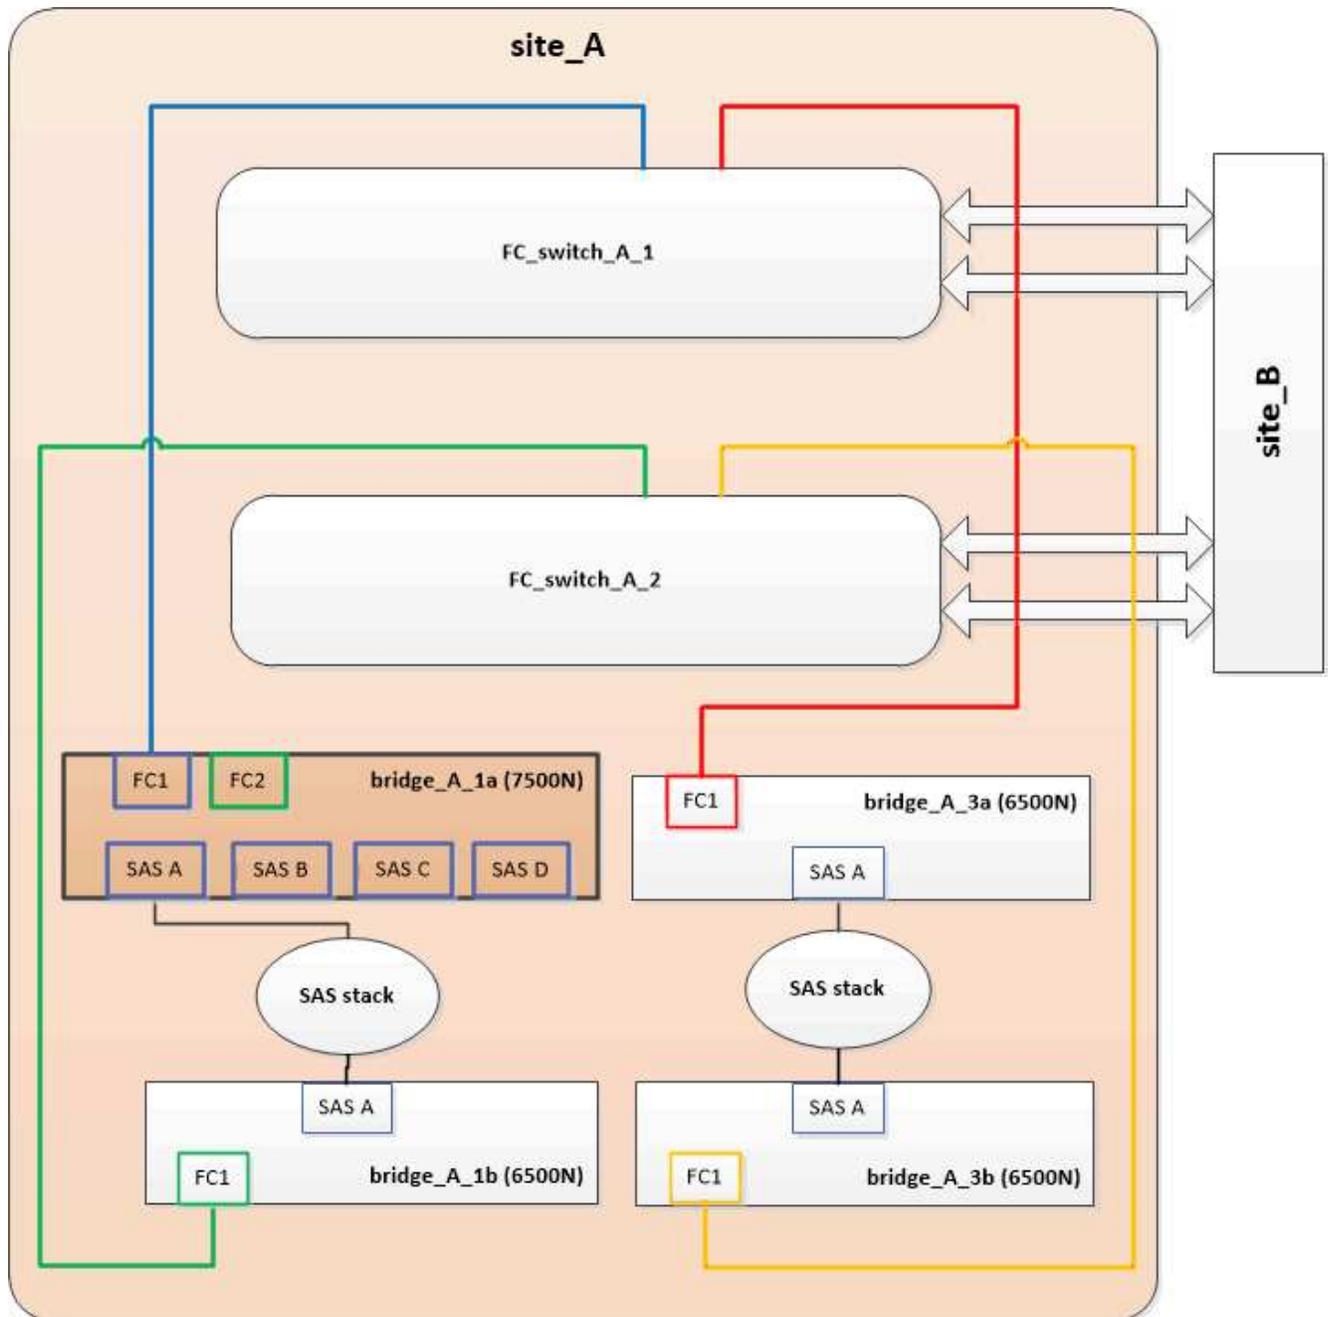
The following diagram shows an example of the initial configuration, in which four FibreBridge 6500N bridges are connecting two SAS storage stacks:



Steps

1. Using the following guidelines, hot-swap the top FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge using the procedure in [Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#):
 - Connect the FibreBridge 7600N or 7500N bridge FC1 port to the switch or controller.
This is the same connection that was made to the FibreBridge 6500N bridge FC1 port.
 - Do not connect the FibreBridge 7600N or 7500N bridge FC2 port at this time.
The following diagram shows that bridge_A_1a has been replaced and is now a FibreBridge 7600N or

7500N bridge:



2. Confirm connectivity to the bridge-connected disks and that the new FibreBridge 7500N is visible in the configuration:

```
run local sysconfig -v
```

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
```

```

.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
847.5GB (1953525168 512B/sect)
    brcd6505-fcs40:12.126L1527      : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
    brcd6505-fcs40:12.126L1528      : NETAPP  X302_HJUPI01TSSA NA02
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
    brcd6505-fcs40:12.126L0          : ATTO    FibreBridge7500N A30H
FB7500N100104**<===**
    brcd6505-fcs42:13.126L0         : ATTO    FibreBridge6500N 1.61
FB6500N102980
    brcd6505-fcs42:6.126L0          : ATTO    FibreBridge6500N 1.61
FB6500N101167
    brcd6505-fcs42:7.126L0         : ATTO    FibreBridge6500N 1.61
FB6500N102974
.
.
.
**<List of storage shelves visible to port\>**
    brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
    brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.
.
.

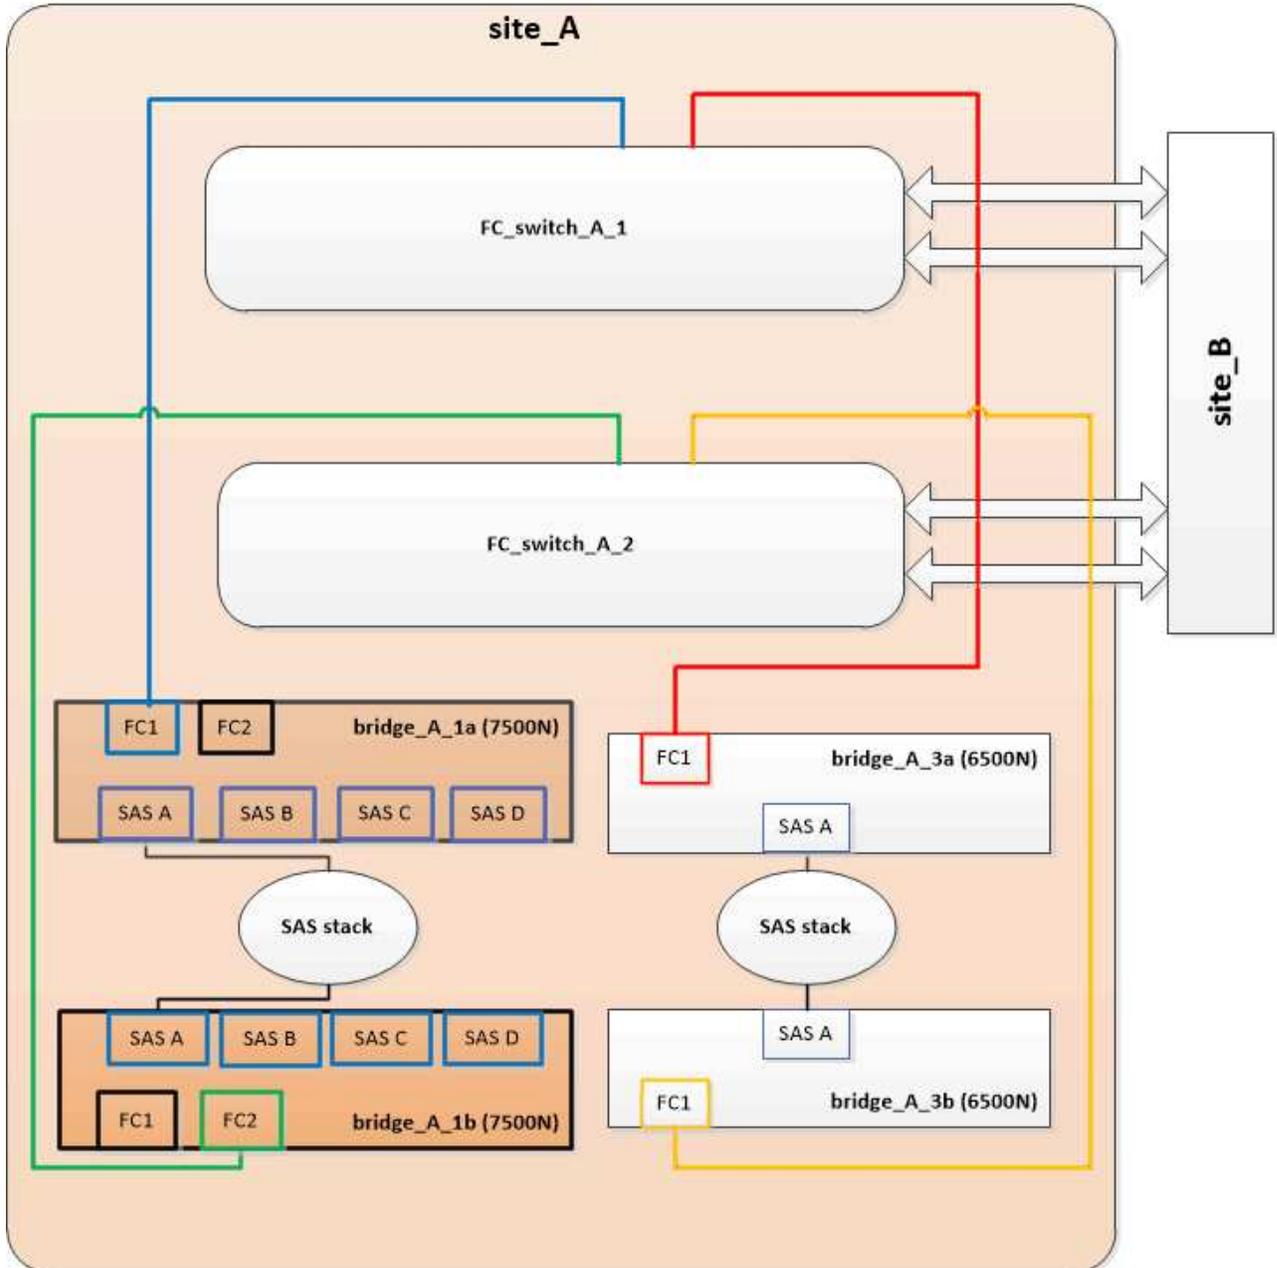
```

3. Using the following guidelines, hot-swap the bottom FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge using the procedure in [Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#):

- Connect the FibreBridge 7600N or 7500N bridge FC2 port to the switch or controller.

This is the same connection that was made to the FibreBridge 6500N bridge FC1 port.

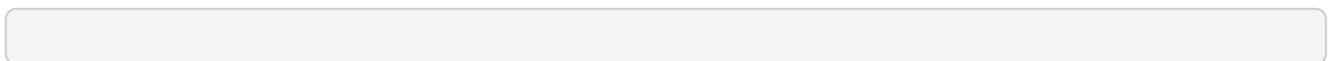
- Do not connect the FibreBridge 7600N or 7500N bridge FC1 port at this time.



4. Confirm connectivity to the bridge-connected disks:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:



```

node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
brcd6505-fcs40:12.126L1527 : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528 : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0 : ATTO      FibreBridge7500N A30H
FB7500N100104
brcd6505-fcs42:13.126L0 : ATTO      FibreBridge7500N A30H
FB7500N100104
.
.
.
**<List of storage shelves visible to port\>**
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200

```

•
•
•

Cabling the bridge SAS ports when consolidating storage behind FibreBridge 7600N or 7500N bridges

When consolidating multiple SAS storage stacks behind a single pair of FibreBridge 7600N or 7500N bridges with available SAS ports, you must move the top and bottom SAS cables to the new bridges.

About this task

The FibreBridge 6500N bridge SAS ports use QSFP connectors. The FibreBridge 7600N or 7500N bridge SAS ports use mini-SAS connectors.



If you insert a SAS cable into the wrong port, when you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that the cable has been moved to another port.

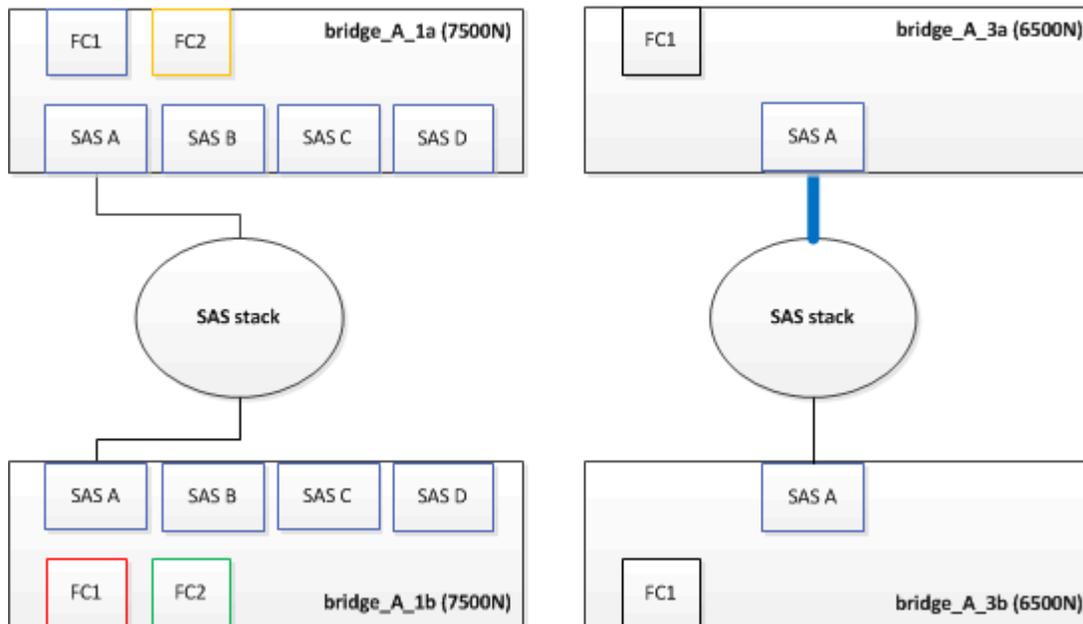


Wait at least 10 seconds before connecting the port. The SAS cable connectors are keyed; when oriented correctly into a SAS port, the connector clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented down (on the underside of the connector).

Steps

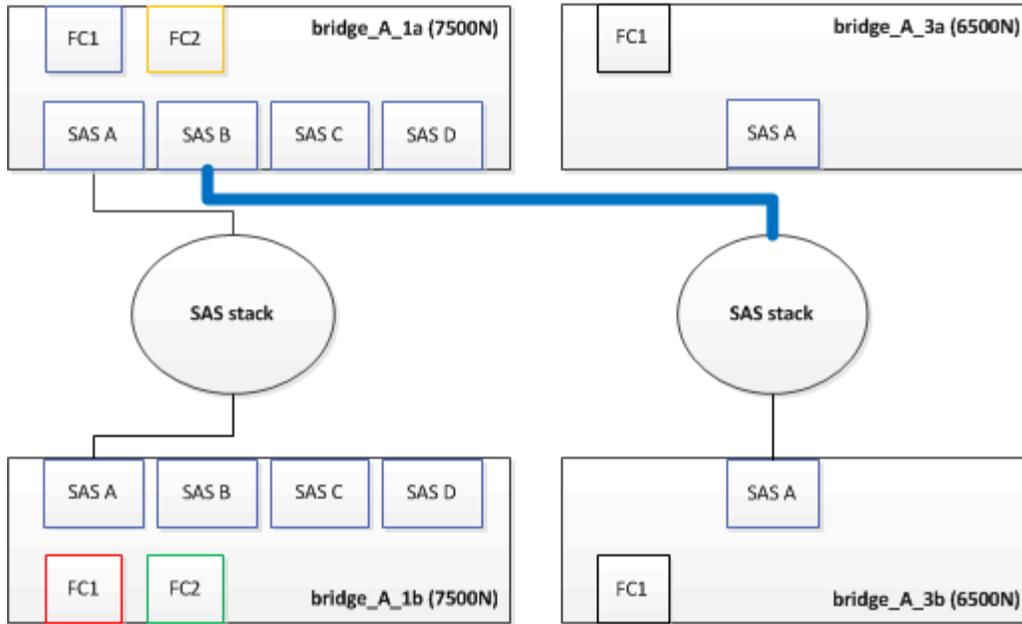
1. Remove the cable that connects the SAS A port of the top FibreBridge 6500N bridge to the top SAS shelf, being sure to note the SAS port on the storage shelf to which it connects.

The cable is shown in blue in the following example:



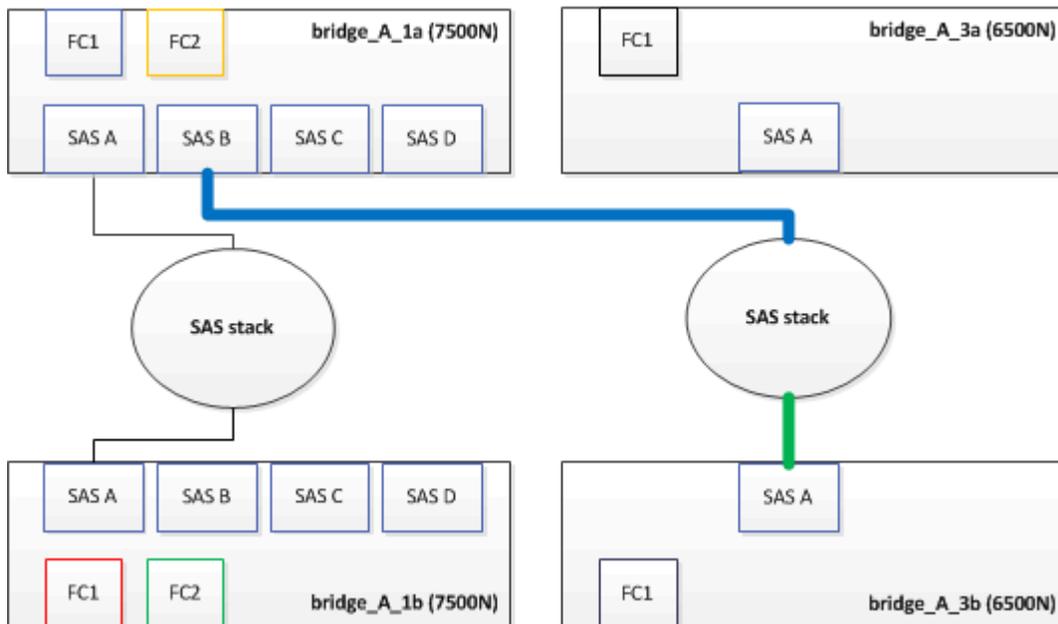
2. Using a cable with a mini-SAS connector, connect the same SAS port on the storage shelf to the SAS B port of the top FibreBridge 7600N or 7500N bridge.

The cable is shown in blue in the following example:



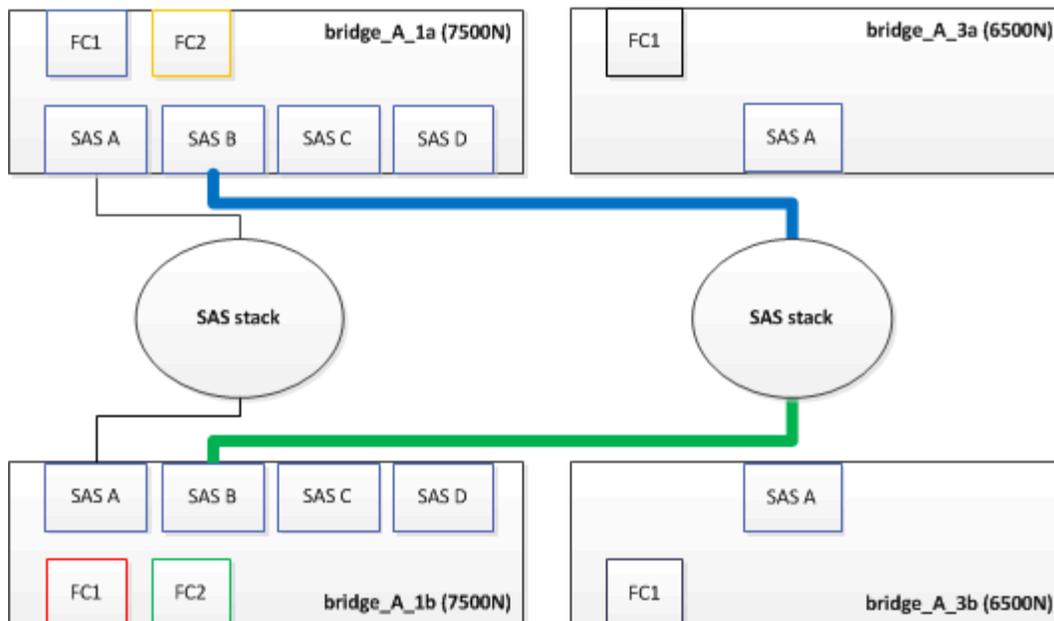
3. Remove the cable that connects the SAS A port of the bottom FibreBridge 6500N bridge to the top SAS shelf, being sure to note the SAS port on the storage shelf to which it connects.

This cable is shown in green in the following example:



4. Using a cable with a mini-SAS connector, connect the same SAS port on the storage shelf to the SAS B port of the bottom FibreBridge 7600N or 7500N bridge.

This cable is shown in green in the following example:



5. Confirm connectivity to the bridge-connected disks:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
```

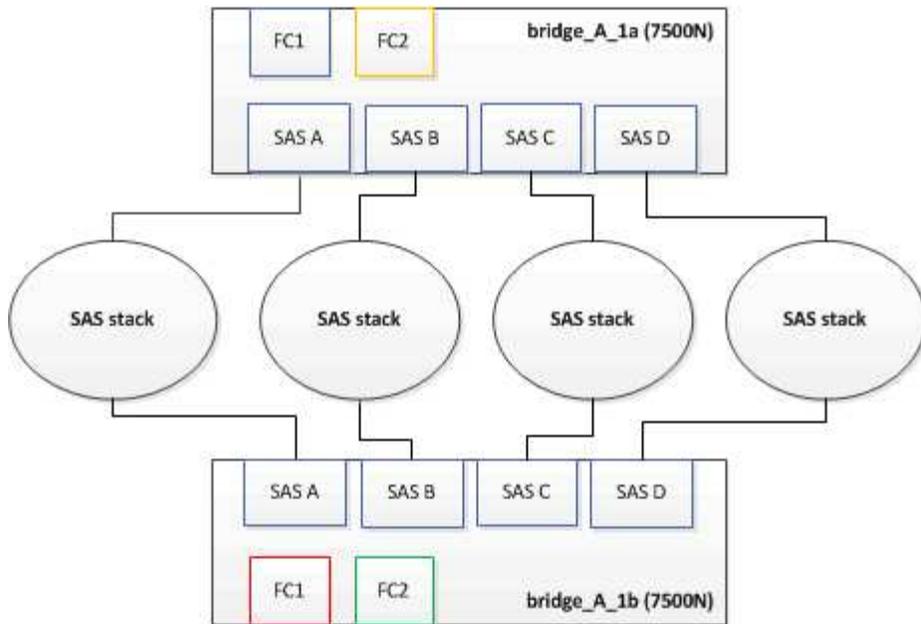
```

**<List of disks visible to port\>**
      ID      Vendor   Model                               FW      Size
      brcd6505-fcs40:12.126L1527    : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
      brcd6505-fcs40:12.126L1528    : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
      .
      .
      .
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
      brcd6505-fcs40:12.126L0        : ATTO     FibreBridge7500N A30H
FB7500N100104
      brcd6505-fcs42:13.126L0        : ATTO     FibreBridge7500N A30H
FB7500N100104
      .
      .
      .
**<List of storage shelves visible to port\>**
      brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
      .
      .
      .

```

6. Remove the old FibreBridge 6500N bridges that are no longer connected to the SAS storage.
7. Wait two minutes for the system to recognize the changes.
8. If the system was cabled incorrectly, remove the cable, correct the cabling, and then reconnect the correct cable.
9. If necessary, repeat the preceding steps to move up to two additional SAS stacks behind the new FibreBridge 7600N or 7500N bridges, using SAS ports C and then D.

Each SAS stack must be connected to the same SAS port on the top and bottom bridge. For example, if the top connection of the stack is connected to the top bridge SAS B port, the bottom connection must be connected to the SAS B port of the bottom bridge.



Updating zoning when adding FibreBridge 7600N or 7500N bridges to a configuration

The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7600N or 7500N bridges and using both FC ports on the FibreBridge 7600N or 7500N bridges. The required changes depend on whether you are running a version of ONTAP earlier than 9.1 or 9.1 and later.

Updating zoning when adding FibreBridge 7500N bridges to a configuration (prior to ONTAP 9.1)

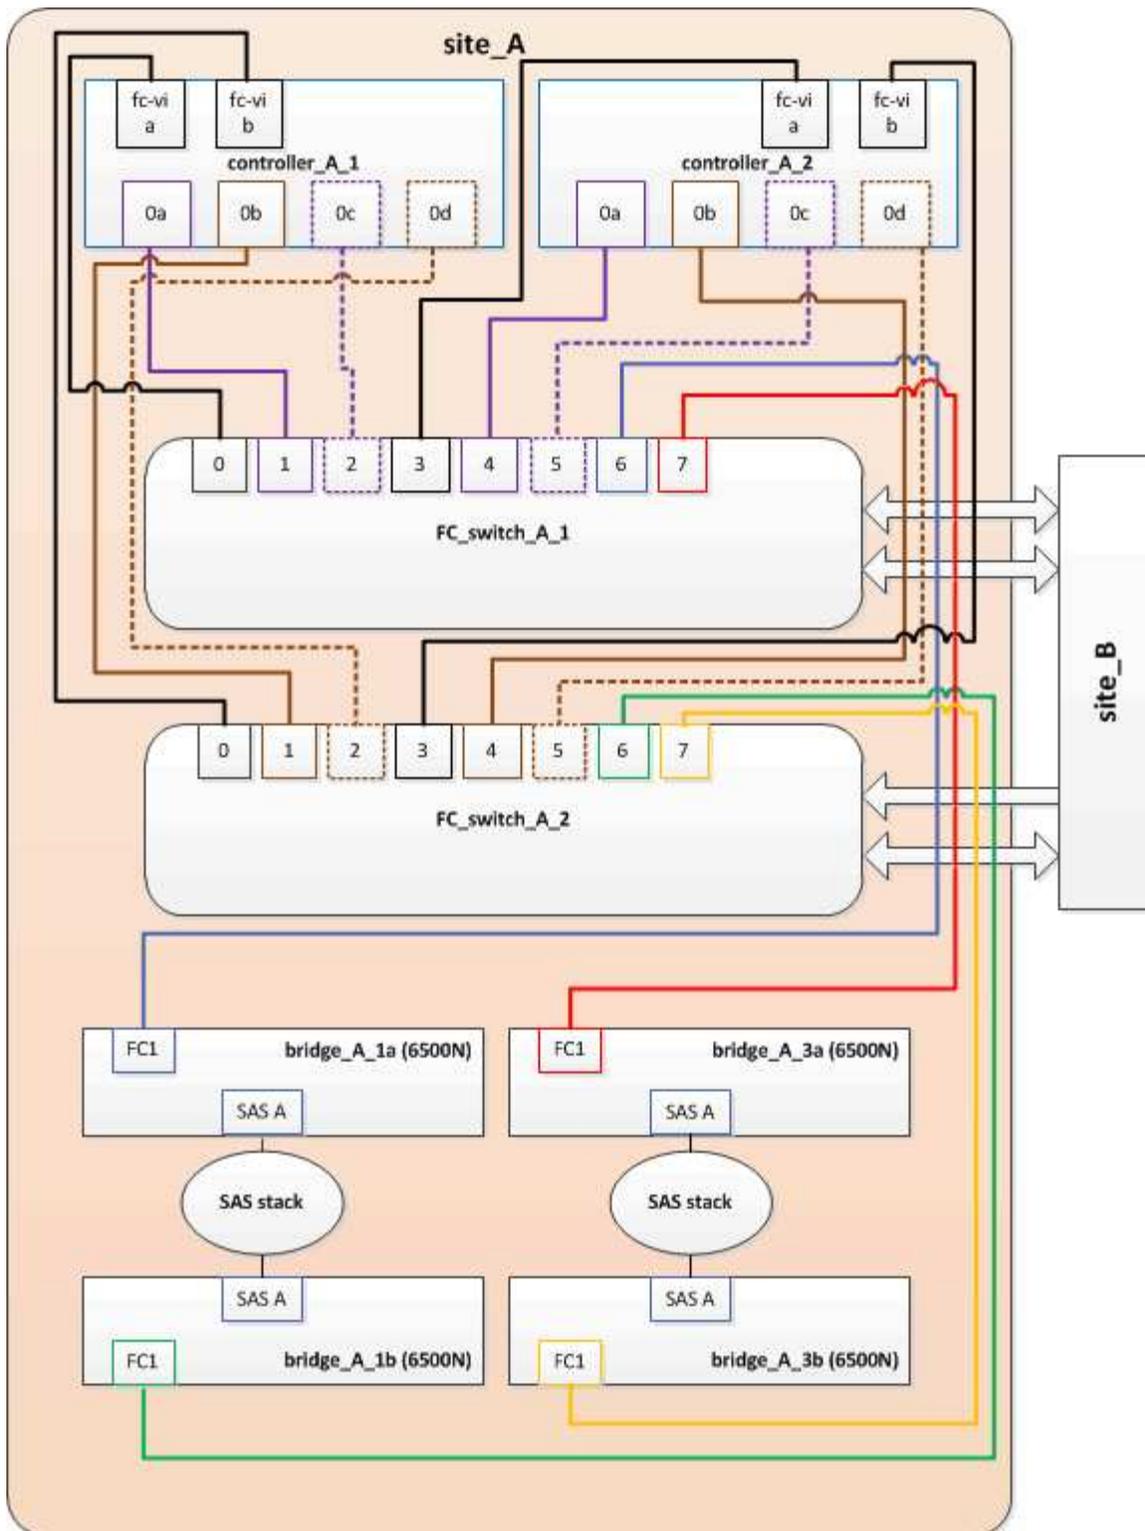
The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7500N bridges and using both FC ports on the FibreBridge 7500N bridges. Each zone can have no more than four initiator ports. The zoning you use depends on whether you are running ONTAP prior to version 9.1 or 9.1 and later.

About this task

The specific zoning in this task is for versions of ONTAP prior to version 9.1.

The zoning changes are required to avoid issues with ONTAP, which requires that no more than four FC initiator ports can have a path to a disk. After recabling to consolidate the shelves, the existing zoning would result in each disk being reachable by eight FC ports. You must change the zoning to reduce the initiator ports in each zone to four.

The following diagram shows the zoning on site_A before the changes:



Steps

1. Update the storage zones for the FC switches by removing half of the initiator ports from each existing zone and creating new zones for the FibreBridge 7500N FC2 ports.

The zones for the new FC2 ports will contain the initiator ports removed from the existing zones. In the diagrams, these zones are shown with dashed lines.

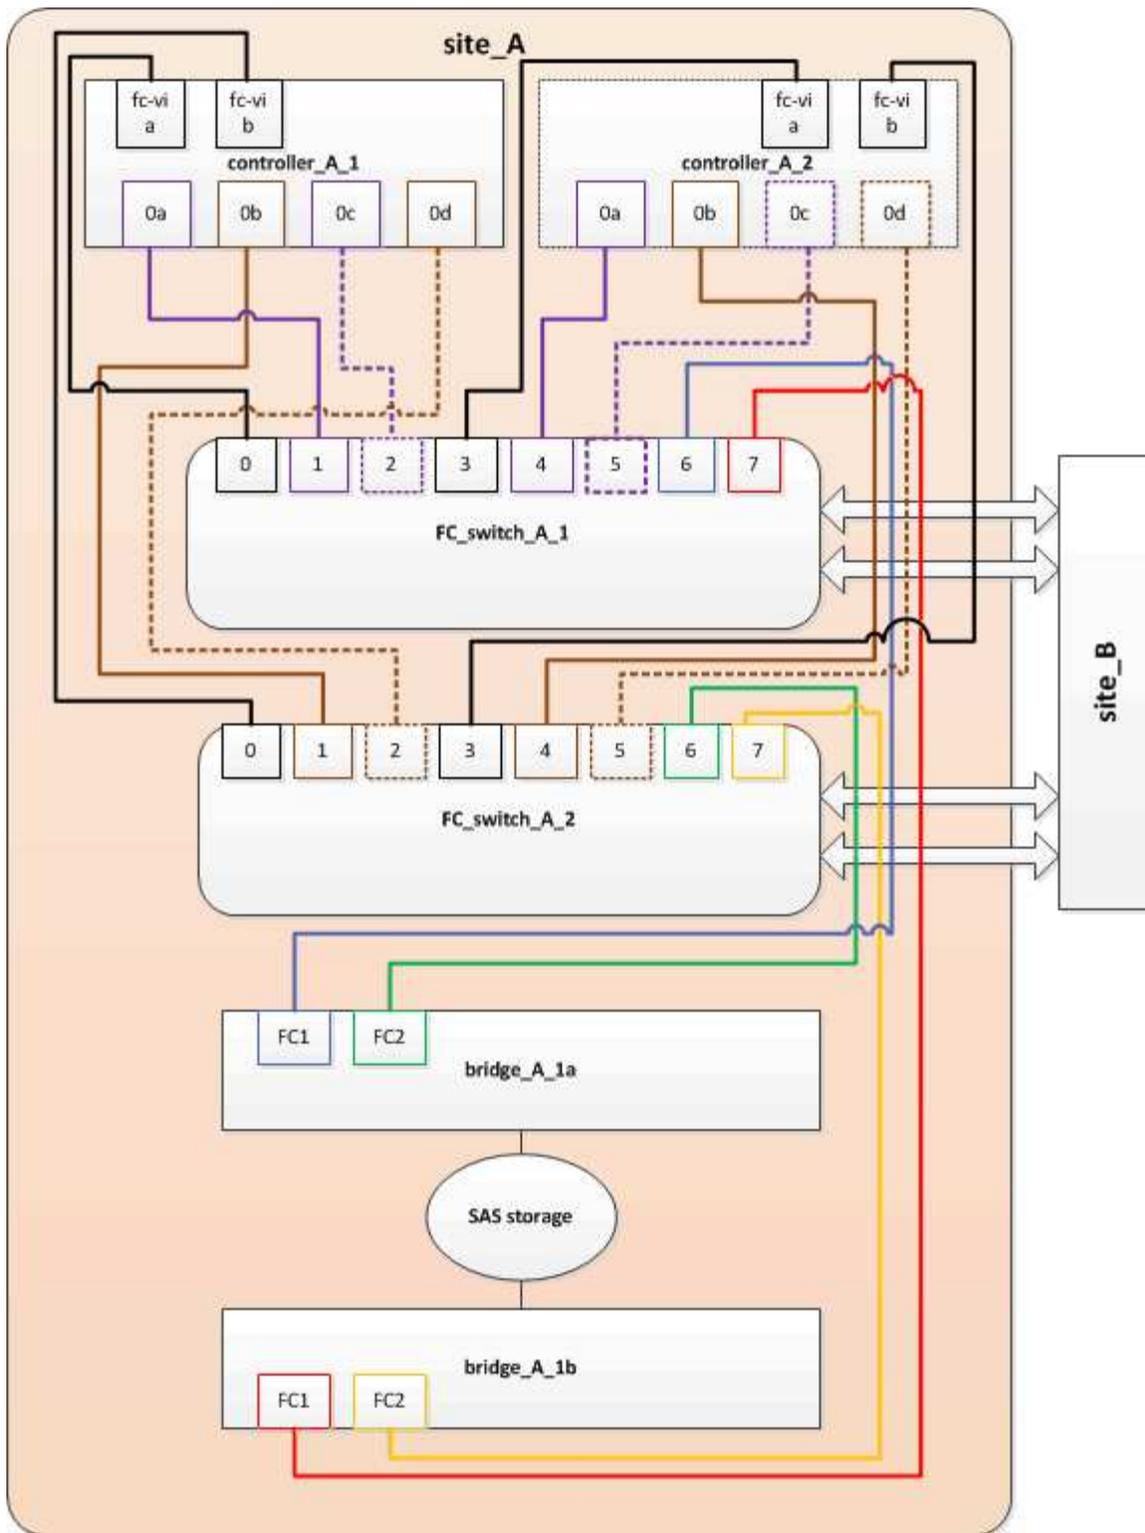
For details about the zoning commands, see the FC switch sections of the [Fabric-attached MetroCluster installation and configuration](#) or [Stretch MetroCluster installation and configuration](#).

The following examples show the storage zones and the ports in each zone before and after the consolidation. The ports are identified by *domain, port* pairs.

- Domain 5 consists of switch FC_switch_A_1.
- Domain 6 consists of switch FC_switch_A_2.
- Domain 7 consists of switch FC_switch_B_1.
- Domain 8 consists of switch FC_switch_B_2.

Before or after consolidation	Zone	Domains and ports	Colors in diagrams (The diagrams only show Site A)
Zones before the consolidation. There is a zone for each FC port on the four FibreBridge 6500N bridges.	STOR_A_1a-FC1	5,1; 5,2; 5,4; 5,5; 7,1; 7,2; 7,4; 7,5; 5,6	Purple + dashed purple + blue
	STOR_A_1b-FC1	6,1; 6,2; 6,4; 6,5; 8,1; 8,2; 8,4; 8,5; 6,6	Brown + dashed brown + green
	STOR_A_2a-FC1	5,1; 5,2; 5,4; 5,5; 7,1; 7,2; 7,4; 7,5; 5,7	Purple + dashed purple + red
	STOR_A_2b-FC1	6,1; 6,2; 6,4; 6,5; 8,1; 8,2; 8,4; 8,5; 6,7	Brown + dashed brown + orange
Zones after the consolidation. There is a zone for each FC port on the two FibreBridge 7500N bridges.	STOR_A_1a-FC1	7,1; 7,4; 5,1; 5,4; 5,6	Purple + blue
	STOR_A_1b-FC1	7,2; 7,5; 5,2; 5,5; 5,7	Dashed purple + red
	STOR_A_1a-FC2	8,1; 8,4; 6,1; 6,4; 6,6	Brown + green
	STOR_A_1b-FC2	8,2; 8,5; 6,2; 6,5; 6,7	Dashed brown + orange

The following diagram shows zoning at site_A after the consolidation:



Updating zoning when adding FibreBridge 7600N or 7500N bridges to a configuration (ONTAP 9.1 and later)

The zoning must be changed when you are replacing FibreBridge 6500N bridges with FibreBridge 7600N or 7500N bridges and using both FC ports on the FibreBridge 7600N or 7500N bridges. Each zone can have no more than four initiator ports.

About this task

- This task applies to ONTAP 9.1 and later.
- FibreBridge 7600N bridges are supported in ONTAP 9.6 and later.
- The specific zoning in this task is for ONTAP 9.1 and later.
- The zoning changes are required to avoid issues with ONTAP, which requires that no more than four FC initiator ports can have a path to a disk.

After recabling to consolidate the shelves, the existing zoning would result in each disk being reachable by eight FC ports. You must change the zoning to reduce the initiator ports in each zone to four.

Step

1. Update the storage zones for the FC switches by removing half of the initiator ports from each existing zone and creating new zones for the FibreBridge 7600N or 7500N FC2 ports.

The zones for the new FC2 ports will contain the initiator ports removed from the existing zones.

Refer to the FC switch section of [Fabric-attached MetroCluster installation and configuration](#) for details about the zoning commands.

Cabling the second bridge FC port when adding FibreBridge 7600N or 7500N bridges to a configuration

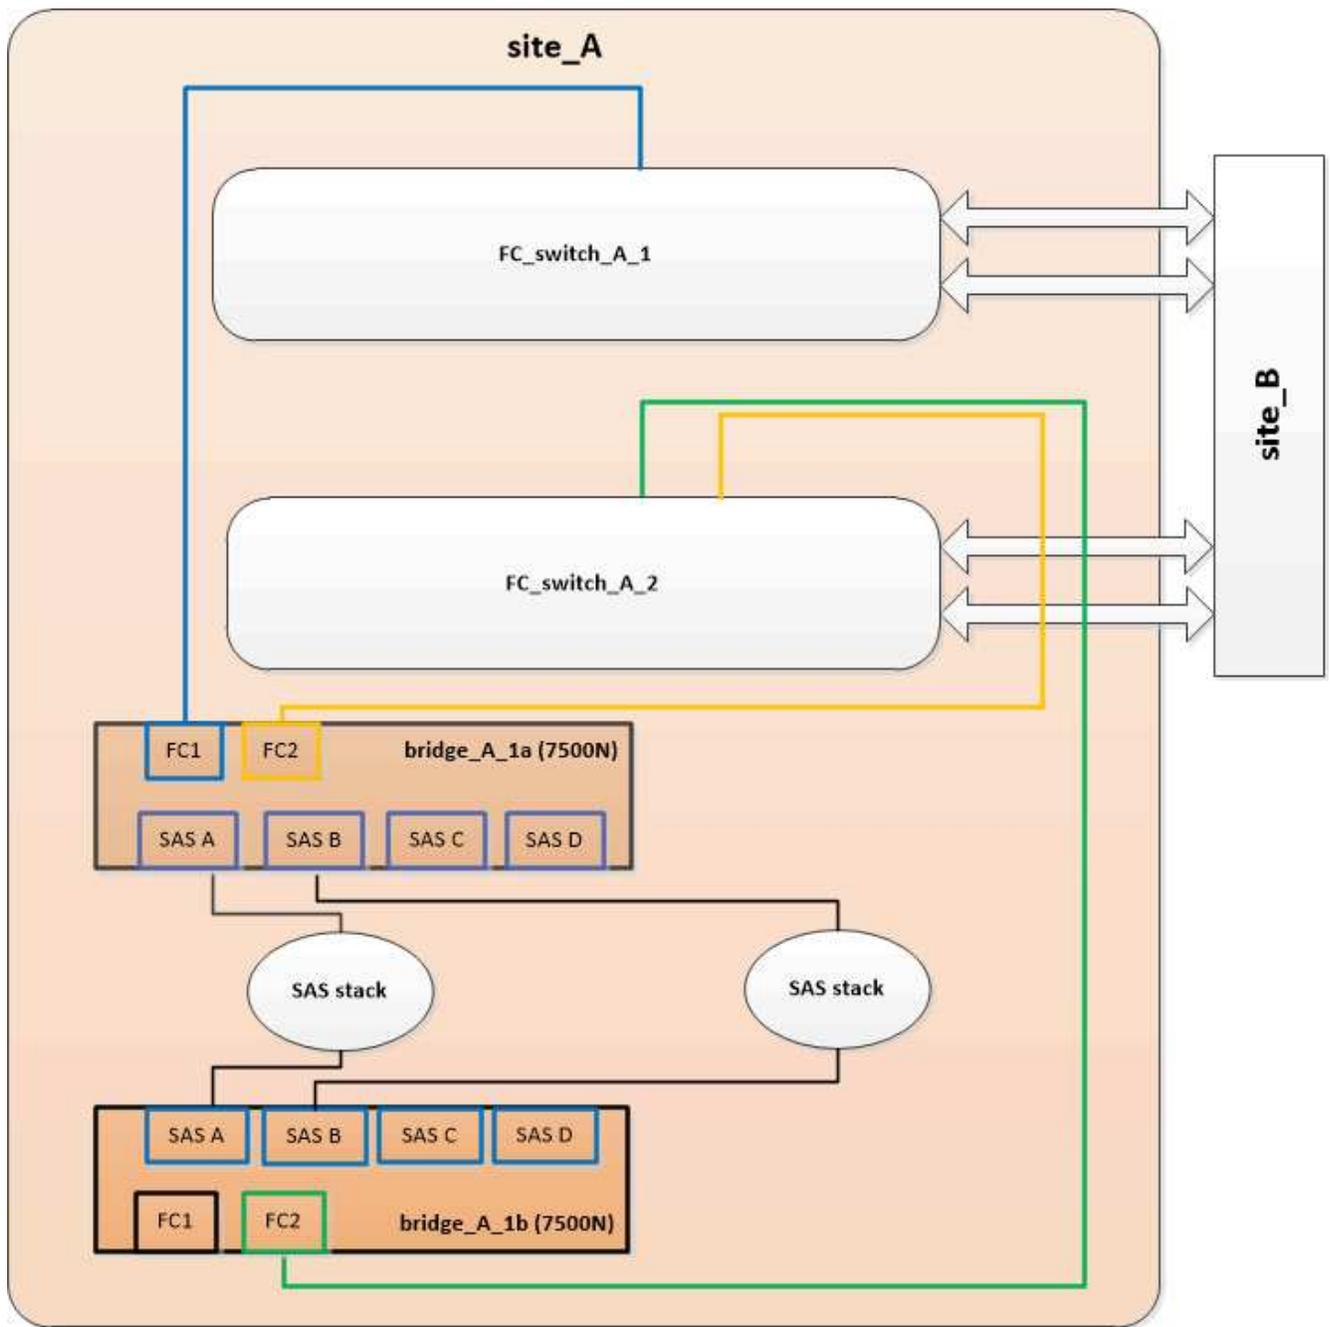
To provide multiple paths to the storage stacks, you can cable the second FC port on each FibreBridge 7600N or 7500N bridge when you have added the FibreBridge 7600N or 7500N bridge to your configuration.

Before you begin

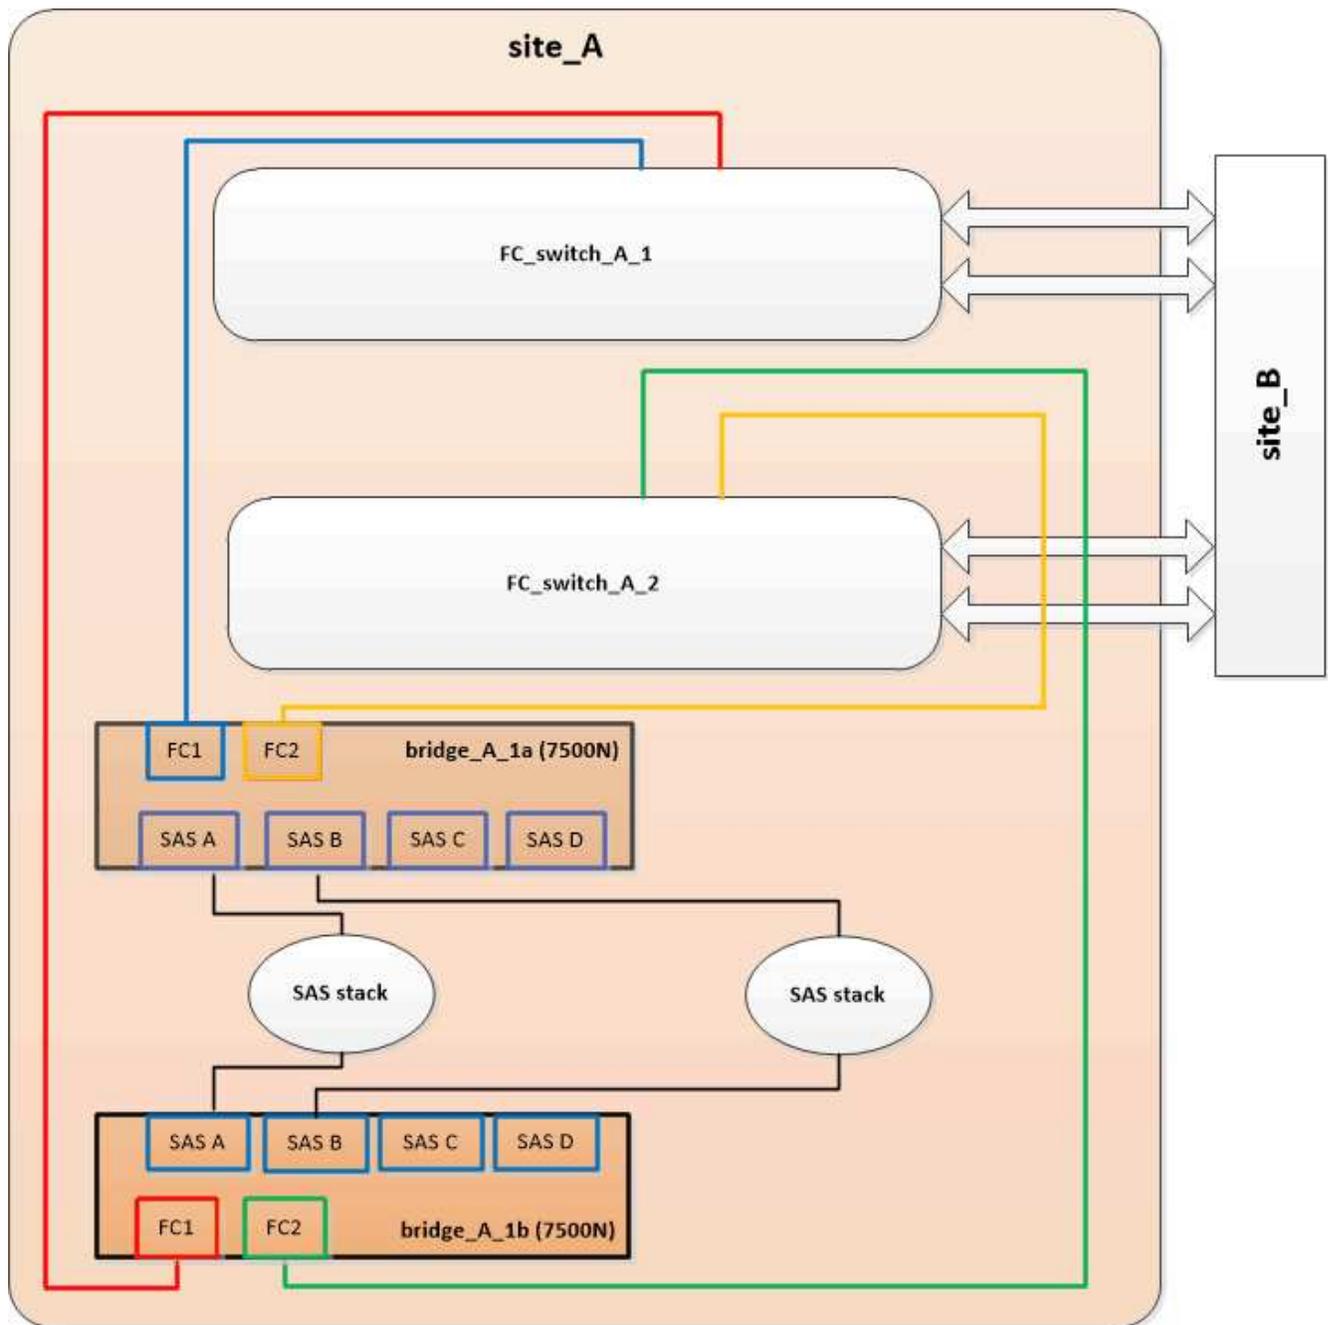
The zoning must have been adjusted to provide zones for the second FC ports.

Steps

1. Cable the FC2 port of the top bridge to the correct port on FC_switch_A_2.



2. Cable the FC1 port of the bottom bridge to the correct port on FC_switch_A_1.



3. Confirm connectivity to the bridge-connected disks:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2015
System ID: 0536872165 (node_A_1); partner ID: 0536872141 (node_B_1)
System Serial Number: 940001025465 (node_A_1)
System Rev: 70
System Storage Configuration: Multi-Path HA**<=== Configuration should
```

```

be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60100
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        FINISAR CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0R1R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
brcd6505-fcs40:12.126L1527 : NETAPP  X302_HJUPI01TSSM NA04
847.5GB (1953525168 512B/sect)
brcd6505-fcs40:12.126L1528 : NETAPP  X302_HJUPI01TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
brcd6505-fcs40:12.126L0      : ATTO      FibreBridge7500N A30H
FB7500N100104
brcd6505-fcs42:13.126L0     : ATTO      FibreBridge7500N A30H
FB7500N100104
.
.
.
**<List of storage shelves visible to port\>**
brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.
.
.

```

Disabling unused SAS ports on the FC-to-SAS bridges

After making cabling changes to the bridge, you should disable any unused SAS ports on FC-to-SAS bridges to avoid health monitor alerts related to the unused ports.

Steps

1. Disable unused SAS ports on the top FC-to-SAS bridge:
 - a. Log in to the bridge CLI.
 - b. Disable any unused ports.



If you have configured an ATTO 7500N bridge, then all of the SAS ports (A through D) are enabled by default, and you must disable the SAS ports that are not being used:

```
SASPortDisable sas port
```

If SAS ports A and B are used, then SAS ports C and D must be disabled. In the following example, the unused SAS ports C and D are disabled:

```
Ready. *
SASPortDisable C

SAS Port C has been disabled.

Ready. *
SASPortDisable D

SAS Port D has been disabled.

Ready. *
```

- c. Save the bridge configuration:

```
SaveConfiguration
```

The following example shows that SAS ports C and D have been disabled. Note that the asterisk no longer appears, indicating that the configuration has been saved.

```
Ready. *
SaveConfiguration

Ready.
```

2. Repeat the previous step on the bottom FC-to-SAS bridge.

Requirements for using other interfaces to configure and manage FibreBridge bridges

You can use the combination of a serial port, Telnet, and FTP to manage the FibreBridge

bridges instead of the recommended management interfaces. Your system must meet the requirements for the applicable interface before you install the bridges.

You can use a serial port or Telnet to configure the bridge and Ethernet management 1 port, and to manage the bridge. You can use FTP to update the bridge firmware.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this document on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

Serial port

When using the serial port to configure and manage a bridge, and to configure the Ethernet management 1 port, your system must meet the following requirements:

- A serial cable (which connects from the bridge serial port to a serial (COM) port on the computer you are using for setup)

The bridge serial port is RJ-45 and has the same pin-out as the controllers.

- A terminal emulation program such as Hyperterminal, Teraterm, or PuTTY to access the console

The terminal program should be capable of logging screen output to a file.

Telnet

When using Telnet to configure and manage a bridge, your system must meet the following requirements:

- A serial cable (which connects from the bridge serial port to a serial (COM) port on the computer you are using for setup)

The bridge serial port is RJ-45 and has the same pin-out as the controllers.

- (Recommended) A non-default user name and password (for accessing the bridge)
- A terminal emulation program such as Hyperterminal, Teraterm, or PuTTY to access the console

The terminal program should be capable of logging screen output to a file.

- An IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge

FTP

When using FTP to update bridge firmware, your system must meet the following requirements:

- A standard Ethernet cable (which connects from the bridge Ethernet management 1 port to your network)
- (Recommended) A non-default user name and password (for accessing the bridge)

Hot-replacing a failed power supply module

When there is a change in status of a power supply module to the bridge, you can remove and install the power supply module.

You can view the change in status of a power supply module through the LEDs on the bridge. You can also view the status of power supply modules via ExpressNAV GUI and the bridge CLI, via serial port, or via Telnet.

- This procedure is NDO (non-disruptive) and takes approximately 15 minutes to complete.
- You need the admin password and access to an FTP or SCP server.



The *ATTO FibreBridge Installation and Operation Manual* for your model bridge has more information about management interfaces.

You can access this and other content on the ATTO web site by using the link provided on the ATTO FibreBridge Description page.

In-band management of the FC-to-SAS bridges

Beginning with ONTAP 9.5 with FibreBridge 7500N or 7600N bridges, in-band management of the bridges is supported as an alternative to IP management of the bridges. Beginning with ONTAP 9.8, out-of-band management is deprecated.



About this task

Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

When using in-band management, the bridges can be managed and monitored from the ONTAP CLI via the FC connection to the bridge. Physical access to the bridge via the bridge Ethernet ports is not required, reducing the security vulnerability of the bridge.

The availability of in-band management of the bridges depends on the version of ONTAP:

- Beginning with ONTAP 9.8, bridges are managed via in-band connections by default and out-of-band management of the bridges via SNMP is deprecated.
- ONTAP 9.5 through 9.7: Either in-band management or out-of-band SNMP management is supported.
- Prior to ONTAP 9.5, only out-of-band SNMP management is supported.

Bridge CLI commands can be issued from the ONTAP interface `storage bridge run-cli -name bridge-name -command bridge-command-name` command at the ONTAP interface.



Using in-band management with IP access disabled is recommended to improve security by limiting physical connectivity the bridge.

Related information

[Hot-swapping a bridge with a replacement bridge of the same model](#)

[Hot-swapping a FibreBridge 7500N with a 7600N bridge](#)

[Hot-swapping a FibreBridge 6500N bridge with a FibreBridge 7600N or 7500N bridge](#)

[Hot-adding a stack of SAS disk shelves and bridges](#)

Managing a FibreBridge bridge from ONTAP

Beginning with ONTAP 9.5, you can use the ONTAP CLI to pass FibreBridge commands to the bridge and display the results of those commands.

About this task



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

Steps

1. Run the applicable FibreBridge command within the `storage bridge run-cli` command:

```
storage bridge run-cli -name bridge-name -command "command-text"
```

The following command runs the FibreBridge `SASPortDisable` command from the ONTAP prompt to disable SAS port b on the bridge:

```
cluster_A::> storage bridge run-cli -name "SASPortDisable b"

SAS Port B has been disabled.
Ready
cluster_A::>
```

Securing or unsecuring the FibreBridge bridge

To easily disable potentially unsecure Ethernet protocols on a bridge, beginning with ONTAP 9.5 you can secure the bridge. This disables the bridge's Ethernet ports. You can also reenablen Ethernet access.

- Securing the bridge disables telnet and other IP port protocols and services (FTP, ExpressNAV, ICMP, or QuickNAV) on the bridge.
- This procedure uses out-of-band management using the ONTAP prompt, which is available beginning with ONTAP 9.5.

You can issue the commands from the bridge CLI if you are not using out-of-band management.

- The **unsecurebridge** command can be used to reenablen the Ethernet ports.
- In ONTAP 9.7 and earlier, running the **securebridge** command on the ATTO FibreBridge might not update the bridge status correctly on the partner cluster. If this occurs, run the **securebridge** command from the partner cluster.



Beginning with ONTAP 9.8, the **storage bridge** command is replaced with **system bridge**. The following steps show the **storage bridge** command, but if you are running ONTAP 9.8 or later, the **system bridge** command is preferred.

Steps

1. From the ONTAP prompt of the cluster containing the bridge, secure or unsecure the bridge.

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
securebridge
```

The following command unsecures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
unsecurebridge
```

2. From the ONTAP prompt of the cluster containing the bridge, save the bridge configuration:

storage bridge run-cli -bridge *bridge-name* -command saveconfiguration

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
saveconfiguration
```

3. From the ONTAP prompt of the cluster containing the bridge, restart the bridge's firmware:

storage bridge run-cli -bridge *bridge-name* -command firmwarerestart

The following command secures bridge_A_1:

```
cluster_A> storage bridge run-cli -bridge bridge_A_1 -command
firmwarerestart
```

FC switch maintenance and replacement

Upgrade or downgrade the firmware on a Brocade FC switch

To upgrade or downgrade the firmware on a Brocade FC switch, you must use the Brocade-specific commands to disable the switch, perform and verify the firmware change, and reboot and reenale the switch.

About this task

Confirm that you have checked and performed the following tasks for your configuration:

- Verify that your new firmware version is compatible with your solution. See the [Hardware Universe](#) for more information.
- You have the firmware files.

- The system is properly cabled.
- All paths to the storage shelves are available.
- The disk shelf stacks are stable.
- The FC switch fabric is healthy.
- No failed components are present in the system.
- The system is operating normally.
- You have the admin password and access to an FTP or SCP server.
- Console logging is enabled.

[Enable console logging](#)

The switch fabric is disabled during a firmware upgrade or downgrade, and the MetroCluster configuration relies on the second fabric to continue operation.

Beginning in Fabric OS 9.0.1, SNMPv2 is not supported on Brocade switches. If you upgrade to Fabric OS 9.0.1 or later, you must use SNMPv3 for health monitoring. For more information, see [Configuring SNMPv3 in a MetroCluster configuration](#).

If you are upgrading to Fabric OS v 9.2.x or later, you must have a Brocade TruFOS certificate installed, refer to [Brocade Fabric OS Software Upgrade Guide, 9.2.x](#) for more information.

This task must be performed on each of the switch fabrics in succession so that all switches are running the same firmware version.



This procedure is nondisruptive and takes approximately one hour to complete.

Steps

1. Log in to each of the switches in the fabric.

The examples in the following steps use the switch `FC_switch_A_1`.

2. Disable each of the switches in the fabric:

`switchCfgPersistentDisable`

If this command is not available, then run the `switchDisable` command.

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

3. Download the desired firmware version:

`firmwareDownload`

When prompted for the file name, you must specify the subdirectory or relative path to the firmware file.

You can run the `firmwareDownload` command at the same time on both switches, but you must allow the firmware to download and commit properly before moving to the next step.

```
FC_switch_A_1:admin> firmwaredownload
Server Name or IP Address: 10.64.203.188
User Name: test
File Name: v7.3.1b
Network Protocol(1-auto-select, 2-FTP, 3-SCP, 4-SFTP, 5-HTTP) [1]: 2
Password:
Server IP: 10.64.203.188, Protocol IPv4
Checking system settings for firmwaredownload...
System settings check passed.
```

4. Verify that the firmware was downloaded and committed to both partitions:

firmwareShow

The following example shows that the firmware download is complete as both images are updated:

```
FC_switch_A_1:admin> firmwareShow
Appl      Primary/Secondary Versions
-----
FOS       v7.3.1b
          v7.3.1b
```

5. Reboot the switches:

reboot

Some firmware versions automatically perform an haReboot operation after the firmware download is finished. The reboot in this step is required even if the haReboot has been performed.

```
FC_switch_A_1:admin> reboot
```

6. Check whether the new firmware is for an intermediate firmware level or for a final specified release.

If the download is for the intermediate firmware level, then perform the previous two steps until the specified release is installed.

7. Enable the switches:

switchCfgPersistentEnable

If this command is not available, then the switch should be in the `enabled` state after the `reboot` command is executed.

```
FC_switch_A_1:admin> switchCfgPersistentEnable
```

8. Verify that the switches are online and that all of the devices are properly logged in:

switchShow

```
FC_switch_A_1:admin> switchShow
```

9. Verify that the buffer usage information for a port group or all of the port groups in the switch is displayed properly:

portbuffershow

```
FC_switch_A_1:admin> portbuffershow
```

10. Verify that the current configuration of a port is displayed properly:

portcfgshow

```
FC_switch_A_1:admin> portcfgshow
```

Verify the port settings, such as speed, mode, trunking, encryption, and compression, in the Inter-Switch Link (ISL) output. Verify that the port settings were not affected by the firmware download.

11. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. Wait 15 minutes before repeating this procedure for the second switch fabric.

Upgrading or downgrading the firmware on a Cisco FC switch

To upgrade or downgrade the firmware on a Cisco FC switch you must use the Cisco-specific commands to disable the switch, perform and verify the upgrade, and reboot and reenable the switch.

About this task

Confirm that you have checked and performed the following tasks for your configuration:

- The system is properly cabled.
- All paths to the storage shelves are available.
- The disk shelf stacks are stable.
- The FC switch fabric are healthy.
- All components in the system are healthy.
- The system is operating normally.
- You have the admin password and access to an FTP or SCP server.
- Console logging is enabled.

[Enable console logging](#)

The switch fabric is disabled during the firmware upgrade or downgrade and the MetroCluster configuration relies on the second fabric to continue operation.

You must repeat this task on each of the switch fabrics in succession to ensure that all switches are running the same firmware version.

You must have the firmware files.



This procedure is nondisruptive and takes approximately one hour to complete.

Steps

1. Log in to each of the switches in the fabric.

In the examples, the switches are called FC_switch_A_1 and FC_switch_B_1.

2. Determine whether there is enough space in the bootflash directory on each switch:

```
dir bootflash
```

If not, delete the unwanted firmware files by using the `delete bootflash:file_name` command.

3. Copy the kickstart and system files to the switches:

copy source_file target_file

In the following example, the kickstart file (m9200-s2ek9-kickstart-mz.5.2.1.bin) and the system file (m9200-s2ek9-mz.5.2.1.bin) are located on the FTP server 10.10.10.55 in the /firmware/ path.

The following example shows the commands issued on FC_switch_A_1:

```
FC_switch_A_1# copy ftp://10.10.10.55/firmware/m9200-s2ek9-kickstart-  
mz.5.2.1.bin bootflash:m9200-s2ek9-kickstart-mz.5.2.1.bin  
FC_switch_A_1# copy ftp://10.10.10.55/firmware/m9200-s2ek9-mz.5.2.1.bin  
bootflash:m9200-s2ek9-mz.5.2.1.bin
```

4. Disable all of the VSANs on both of the switches in this fabric.

Use the following procedure to disable the VSANs:

- a. Open the config terminal:

config t

- b. Enter: **vsan database**

- c. Check the state of the VSANs:

show vsan

All VSANs must be active.

- d. Suspend the VSANs:

vsan vsan-num suspend

Example: vsan 10 suspend

- e. Check the state of the VSANs again:

show vsan

All VSANs must be suspended.

- f. Exit the config terminal:

end

- g. Save the configuration.

copy running-config startup-config

The following example displays the output for FC_switch_A_1:

```

FC_switch_A_1# config t
Enter configuration commands, one per line.  End with CNTL/Z.
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# show vsan
vsan 1 information
    name:VSAN0001  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:up

vsan 30 information
    name:MC1_FCVI_2_30  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:up

vsan 40 information
    name:MC1_STOR_2_40  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:up

vsan 70 information
    name:MC2_FCVI_2_70  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:up

vsan 80 information
    name:MC2_STOR_2_80  state:active
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:up

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

FC_switch_A_1(config-vsan-db)# vsan 1 suspend
FC_switch_A_1(config-vsan-db)# vsan 30 suspend
FC_switch_A_1(config-vsan-db)# vsan 40 suspend
FC_switch_A_1(config-vsan-db)# vsan 70 suspend
FC_switch_A_1(config-vsan-db)# vsan 80 suspend
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1#
FC_switch_A_1# show vsan

```

```

vsan 1 information
    name:VSAN0001  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:down

vsan 30 information
    name:MC1_FCVI_2_30  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:down

vsan 40 information
    name:MC1_STOR_2_40  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:down

vsan 70 information
    name:MC2_FCVI_2_70  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id
    operational state:down

vsan 80 information
    name:MC2_STOR_2_80  state:suspended
    interoperability mode:default
    loadbalancing:src-id/dst-id/oxid
    operational state:down

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

```

5. Install the desired firmware on the switches:

```

install all system bootflash:systemfile_name kickstart
bootflash:kickstartfile_name

```

The following example shows the commands issued on FC_switch_A_1:

```

FC_switch_A_1# install all system bootflash:m9200-s2ek9-mz.5.2.1.bin
kickstart bootflash:m9200-s2ek9-kickstart-mz.5.2.1.bin
Enter Yes to confirm the installation.

```

6. Check the version of the firmware on each switch to make sure the correct version was installed:

show version

7. Enable all of the VSANs on both of the switches in this fabric.

Use the following procedure to enable the VSANs:

- a. Open the config terminal:

```
config t
```

- b. Enter: **vsan database**

- c. Check the state of the VSANs:

```
show vsan
```

The VSANs must be suspended.

- d. Activate the VSANs:

```
no vsan vsan-num suspend
```

Example: no vsan 10 suspend

- e. Check the state of the VSANs again:

```
show vsan
```

All VSANs must be active.

- f. Exit the config terminal:

```
end
```

- g. Save the configuration:

```
copy running-config startup-config
```

The following example displays the output for FC_switch_A_1:

```
FC_switch_A_1# config t
Enter configuration commands, one per line.  End with CNTL/Z.
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# show vsan
vsan 1 information
      name:VSAN0001  state:suspended
      interoperability mode:default
      loadbalancing:src-id/dst-id/oxid
      operational state:down

vsan 30 information
      name:MC1_FCVI_2_30  state:suspended
```

```

interoperability mode:default
loadbalancing:src-id/dst-id
operational state:down

vsan 40 information
  name:MC1_STOR_2_40  state:suspended
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:down

vsan 70 information
  name:MC2_FCVI_2_70  state:suspended
  interoperability mode:default
  loadbalancing:src-id/dst-id
  operational state:down

vsan 80 information
  name:MC2_STOR_2_80  state:suspended
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:down

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

FC_switch_A_1(config-vsan-db)# no vsan 1 suspend
FC_switch_A_1(config-vsan-db)# no vsan 30 suspend
FC_switch_A_1(config-vsan-db)# no vsan 40 suspend
FC_switch_A_1(config-vsan-db)# no vsan 70 suspend
FC_switch_A_1(config-vsan-db)# no vsan 80 suspend
FC_switch_A_1(config-vsan-db)#
FC_switch_A_1(config-vsan-db)# show vsan
vsan 1 information
  name:VSAN0001  state:active
  interoperability mode:default
  loadbalancing:src-id/dst-id/oxid
  operational state:up

vsan 30 information
  name:MC1_FCVI_2_30  state:active
  interoperability mode:default
  loadbalancing:src-id/dst-id
  operational state:up

vsan 40 information

```

```

name:MC1_STOR_2_40 state:active
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:up

vsan 70 information
name:MC2_FCVI_2_70 state:active
interoperability mode:default
loadbalancing:src-id/dst-id
operational state:up

vsan 80 information
name:MC2_STOR_2_80 state:active
interoperability mode:default
loadbalancing:src-id/dst-id/oxid
operational state:up

vsan 4079:evfp_isolated_vsan

vsan 4094:isolated_vsan

FC_switch_A_1(config-vsan-db) # end
FC_switch_A_1#

```

8. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

9. Repeat this procedure for the second switch fabric.

Upgrading to new Brocade FC switches

If you are upgrading to new Brocade FC switches, you must replace the switches in the first fabric, verify that the MetroCluster configuration is fully operational, and then replace the switches in the second fabric.

- The MetroCluster configuration must be healthy and in normal operation.
- The MetroCluster switch fabrics consist of four Brocade switches.

The illustrations in the following steps show current switches.

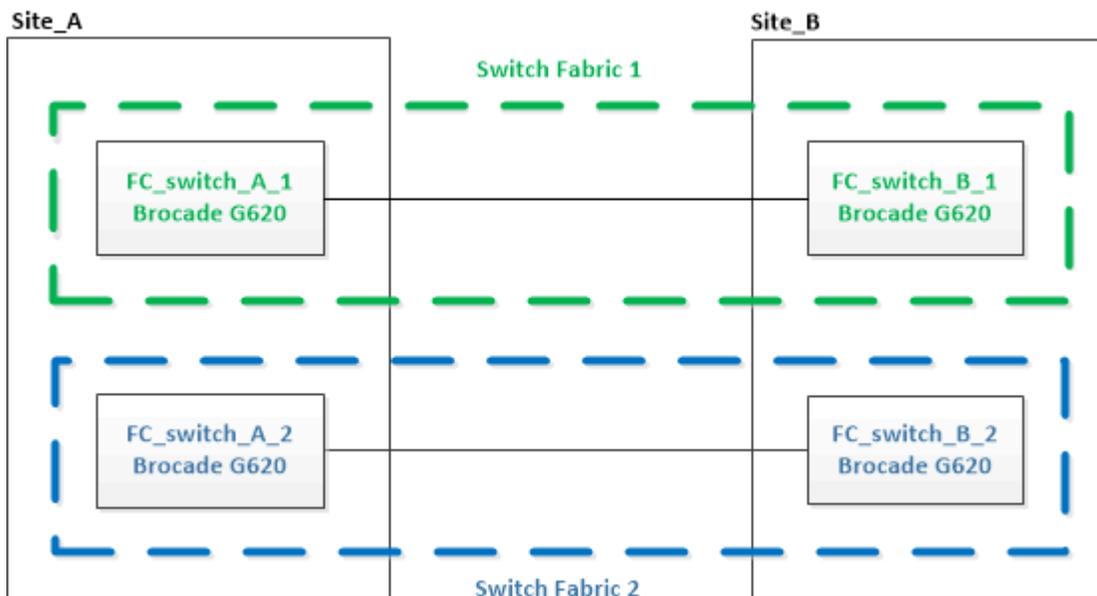
- The switches must be running the most recent supported firmware.

[NetApp Interoperability Matrix Tool](#)

- This procedure is nondisruptive and takes approximately two hours to complete.
- You need the admin password and access to an FTP or SCP server.
- [Enable console logging](#) before performing this task.

The switch fabrics are upgraded one at a time.

At the end of this procedure, all four switches will be upgraded to new switches.

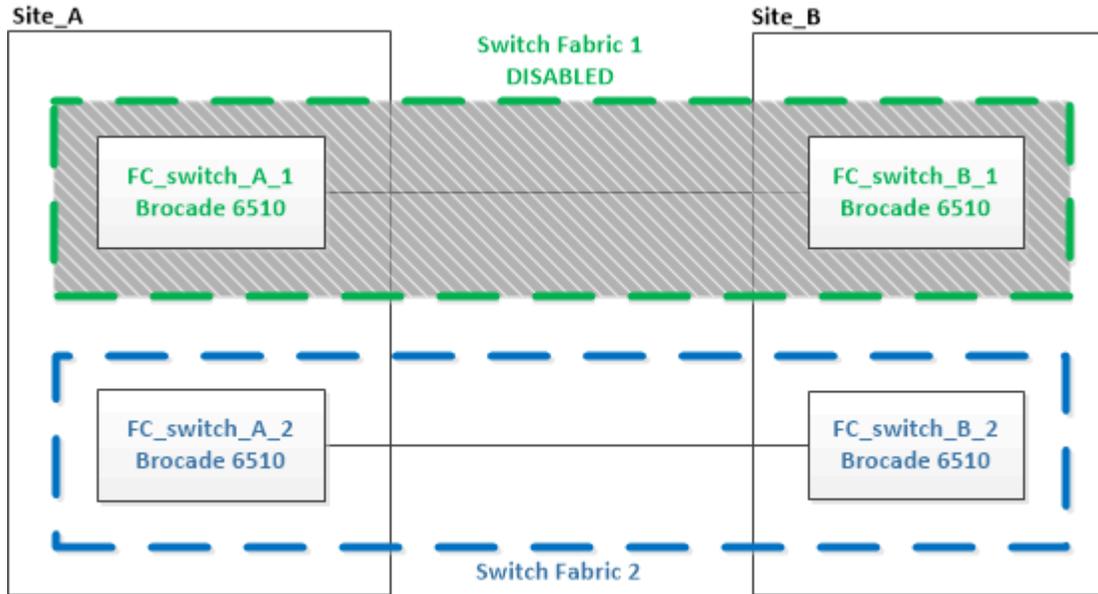


Steps

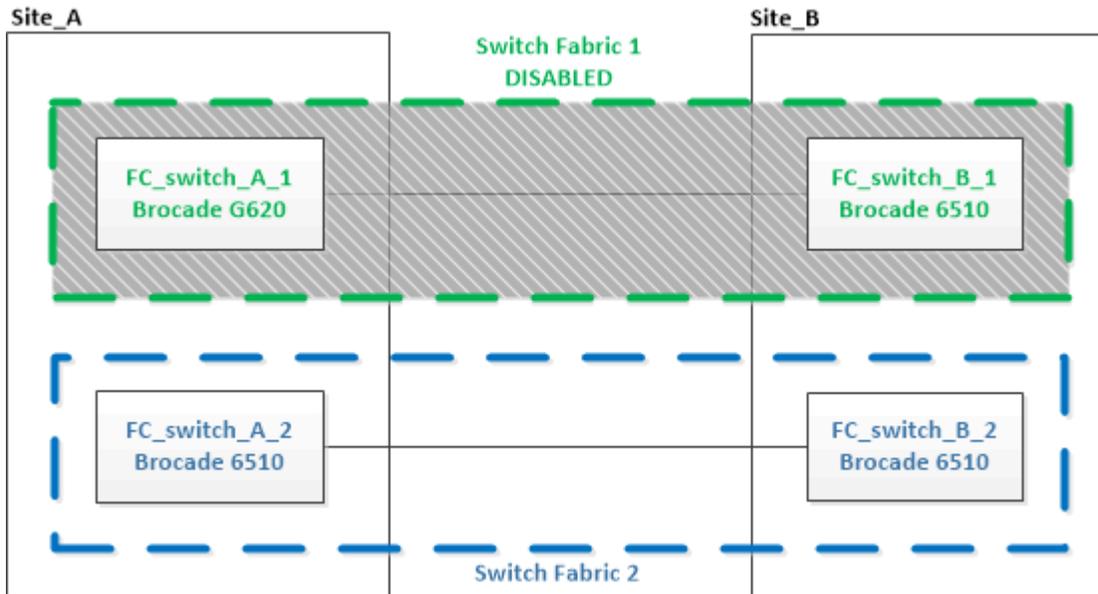
1. Disable the first switch fabric:

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

```
FC_switch_A_1:admin> switchCfgPersistentDisable
```



- 2. Replace the old switches at one MetroCluster site.
 - a. Uncable and remove the disabled switch.
 - b. Install the new switch in the rack.



- c. Disable the new switches by running the following command on both switches:

```
switchCfgPersistentDisable
```

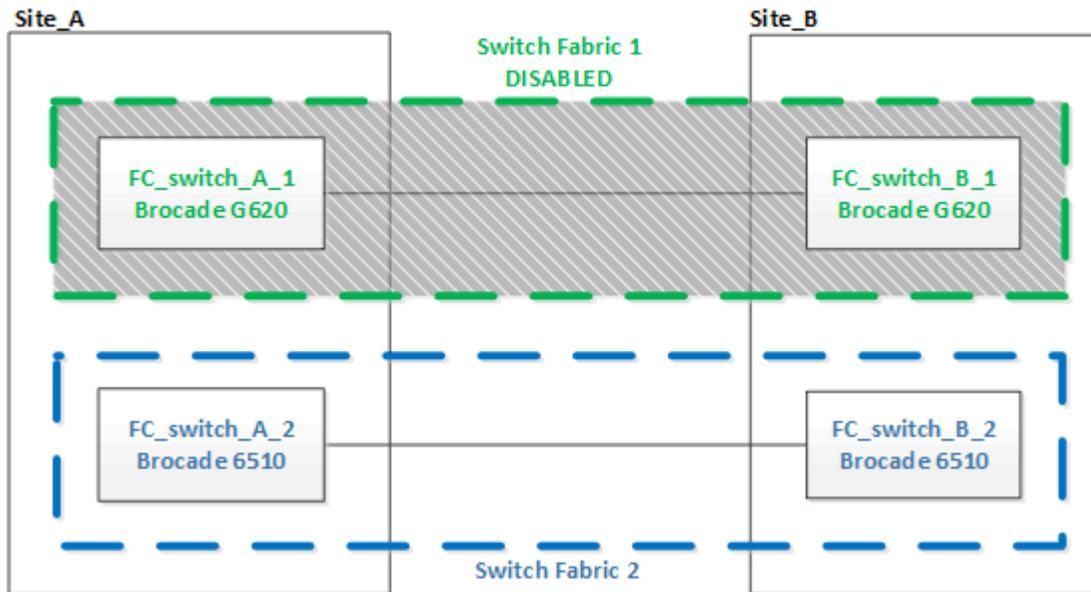
```
FC_switch_A_1:admin> switchCfgPersistentDisable
```

- d. Cable the new switch using the recommended port assignments.

[Port assignments for FC switches](#)

- e. Repeat these substeps at the partner MetroCluster site to replace the second switch in the first switch fabric.

Both switches in fabric 1 have been replaced.



3. Power up the new switches and let them boot up.
4. Configure the Brocade FC switches using one of the following procedures:

[Configure Brocade FC switches with RCF files](#)

[Configure the Brocade FC switches manually](#)

5. Save the switch configuration:

```
cfgSave
```

6. Wait 10 minutes to allow the configuration to stabilize.
7. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

The output shows the disks attached to the initiator ports on the controller, and identifies the shelves connected to the FC-to-SAS bridges:

```
node_A_1> run local sysconfig -v
NetApp Release 9.3.2X18: Sun Dec 13 01:23:24 PST 2017
System ID: 4068741258 (node_A_1); partner ID: 4068741260 (node_B_1)
System Serial Number: 940001025471 (node_A_1)
System Rev: 70
```

```

System Storage Configuration: Multi-Path HA**<=== Configuration should
be multi-path HA**
.
.
.
slot 0: FC Host Adapter 0g (QLogic 8324 rev. 2, N-port, <UP>)**<===
Initiator port**
    Firmware rev:      7.5.0
    Flash rev:         0.0.0
    Host Port Id:      0x60130
    FC Node Name:      5:00a:098201:bae312
    FC Port Name:      5:00a:098201:bae312
    SFP Vendor:        UTILITIES CORP.
    SFP Part Number:   FTLF8529P3BCVAN1
    SFP Serial Number: URQ0Q9R
    SFP Capabilities:  4, 8 or 16 Gbit
    Link Data Rate:    16 Gbit
    Switch Port:       brcd6505-fcs40:1
**<List of disks visible to port\>**
    ID      Vendor  Model          FW      Size
    brcd6505-fcs29:12.126L1527  : NETAPP  X302_HJUPIO1TSSM NA04
847.5GB (1953525168 512B/sect)
    brcd6505-fcs29:12.126L1528  : NETAPP  X302_HJUPIO1TSSA NA02
847.5GB (1953525168 512B/sect)
.
.
.
**<List of FC-to-SAS bridges visible to port\>**
FC-to-SAS Bridge:
    brcd6505-fcs40:12.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
    brcd6505-fcs42:13.126L0      : ATTO      FibreBridge6500N 1.61
FB6500N102980
    brcd6505-fcs42:6.126L0       : ATTO      FibreBridge6500N 1.61
FB6500N101167
    brcd6505-fcs42:7.126L0       : ATTO      FibreBridge6500N 1.61
FB6500N102974
.
.
.
**<List of storage shelves visible to port\>**
    brcd6505-fcs40:12.shelf6: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
    brcd6505-fcs40:12.shelf8: DS4243  Firmware rev. IOM3 A: 0200
IOM3 B: 0200
.

```

·
·

8. Returning to the switch prompt, verify the switch firmware version:

```
firmwareShow
```

The switches must be running the most recent supported firmware.

[NetApp Interoperability Matrix Tool](#)

9. Simulate a switchover operation:

a. From any node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with "y" when prompted to continue into advanced mode and see the advanced mode prompt (*>).

b. Perform the switchover operation with the `-simulate` parameter:

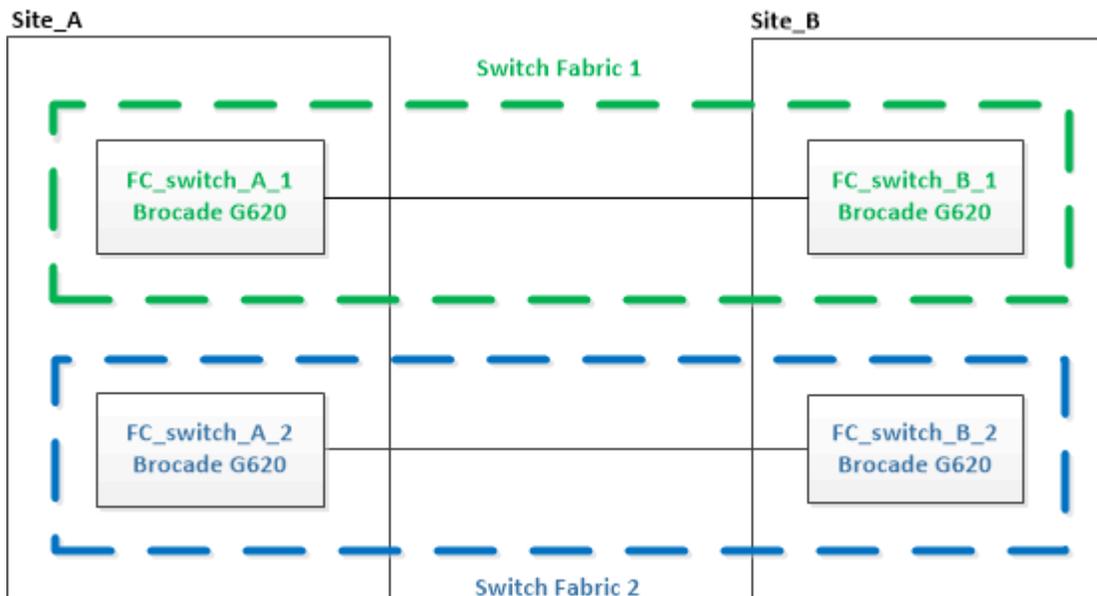
```
metrocluster switchover -simulate
```

c. Return to the admin privilege level:

```
set -privilege admin
```

10. Repeat the previous steps on the second switch fabric.

After repeating the steps, all four switches have been upgraded and the MetroCluster configuration is in normal operation.



Replacing a Brocade FC switch

You must use this Brocade-specific procedure to replace a failed switch.

About this task

You need the admin password and access to an FTP or SCP server.

[Enable console logging](#) before performing this task.

In the following examples, FC_switch_A_1 is the healthy switch and FC_switch_B_1 is the impaired switch. The switch port usage in the examples is shown in the following table:

Port connections	Ports
FC-VI connections	0, 3
HBA connections	1, 2, 4, 5
FC-to-SAS bridge connections	6, 7
ISL connections	10, 11

The examples show two FC-to-SAS bridges. If you have more, you must disable and subsequently enable the additional ports.



This procedure is nondisruptive and takes approximately two hours to complete.

Your switch port usage should follow the recommended assignments.

- [Port assignments for FC switches](#)

Steps

1. Fence off the switch undergoing replacement by disabling the ISL ports on the healthy switch in the fabric and the FC-VI and HBA ports on the impaired switch (if the impaired switch is still operating):
 - a. Disable the ISL ports on the healthy switch for each port:

```
portcfgpersistentdisable port-number
```

```
FC_switch_A_1:admin> portcfgpersistentdisable 10  
FC_switch_A_1:admin> portcfgpersistentdisable 11
```

- b. If the impaired switch is still operational, disable the FC-VI and HBA ports on that switch for each port:

```
portcfgpersistentdisable port-number
```

```
FC_switch_B_1:admin> portcfgpersistentdisable 0
FC_switch_B_1:admin> portcfgpersistentdisable 1
FC_switch_B_1:admin> portcfgpersistentdisable 2
FC_switch_B_1:admin> portcfgpersistentdisable 3
FC_switch_B_1:admin> portcfgpersistentdisable 4
FC_switch_B_1:admin> portcfgpersistentdisable 5
```

2. If the impaired switch is still operational, gather the output from the `switchshow` command.

```
FC_switch_B_1:admin> switchshow
  switchName: FC_switch_B_1
  switchType: 71.2
  switchState: Online
  switchMode: Native
  switchRole: Subordinate
  switchDomain:      2
  switchId:   fffc01
  switchWwn:  10:00:00:05:33:86:89:cb
  zoning:      OFF
  switchBeacon: OFF
```

3. Boot and preconfigure the new switch prior to physically installing it:

- a. Power up the new switch and let it boot up.
- b. Check the firmware version on the switch to confirm that it matches the version of the other FC switches:

```
firmwareShow
```

- c. Configure the new switch by following the Brocade procedures in [Configure the FC switches](#).



At this point, the new switch is not cabled to the MetroCluster configuration.

- d. Disable the FC-VI, HBA, and storage ports on the new switch, and the ports connected to the FC-SAS bridges.

```
FC_switch_B_1:admin> portcfgpersistentdisable 0
FC_switch_B_1:admin> portcfgpersistentdisable 1
FC_switch_B_1:admin> portcfgpersistentdisable 2
FC_switch_B_1:admin> portcfgpersistentdisable 3
FC_switch_B_1:admin> portcfgpersistentdisable 4
FC_switch_B_1:admin> portcfgpersistentdisable 5

FC_switch_B_1:admin> portcfgpersistentdisable 6
FC_switch_B_1:admin> portcfgpersistentdisable 7
```

4. Physically replace the switch:

- a. Power off the impaired FC switch.
- b. Power off the replacement FC switch.
- c. Uncable and remove the impaired switch, carefully noting which cables connected to which ports.
- d. Install the replacement switch in the rack.
- e. Cable the replacement switch exactly as the old switch was cabled.
- f. Power on the new FC switch.

5. To enable ISL encryption, refer to [Configure the Brocade FC switches manually](#).

If you are enabling ISL encryption, you need to complete the following tasks:

- Disable the virtual fabric
- Set the payload
- Set the authentication policy
- Enable ISL encryption on Brocade switches

6. Complete the configuration of the new switch:

- a. Enable the ISLs:

```
portcfgpersistentenable port-number
```

```
FC_switch_B_1:admin> portcfgpersistentenable 10
FC_switch_B_1:admin> portcfgpersistentenable 11
```

- b. Verify the zoning configuration:

```
cfg show
```

- c. On the replacement switch (FC_switch_B_1 in the example), verify that the ISLs are online:

```
switchshow
```

```

FC_switch_B_1:admin> switchshow
switchName: FC_switch_B_1
switchType: 71.2
switchState:Online
switchMode: Native
switchRole: Principal
switchDomain:      4
switchId:   fffc03
switchWwn:  10:00:00:05:33:8c:2e:9a
zoning:     OFF
switchBeacon: OFF

Index Port Address Media Speed State  Proto
=====
...
10  10    030A00 id   16G    Online FC E-Port
10:00:00:05:33:86:89:cb "FC_switch_A_1"
11  11    030B00 id   16G    Online FC E-Port
10:00:00:05:33:86:89:cb "FC_switch_A_1" (downstream)
...

```

- d. Enable the storage ports that connect to the FC bridges.

```

FC_switch_B_1:admin> portcfgpersistentenable 6
FC_switch_B_1:admin> portcfgpersistentenable 7

```

- e. Enable the storage, HBA, and FC-VI ports.

The following example shows the commands used to enable the ports connecting HBA adapters:

```

FC_switch_B_1:admin> portcfgpersistentenable 1
FC_switch_B_1:admin> portcfgpersistentenable 2
FC_switch_B_1:admin> portcfgpersistentenable 4
FC_switch_B_1:admin> portcfgpersistentenable 5

```

The following example shows the commands used to enable the ports connecting the FC-VI adapters:

```

FC_switch_B_1:admin> portcfgpersistentenable 0
FC_switch_B_1:admin> portcfgpersistentenable 3

```

7. Verify that the ports are online:

```
switchshow
```

8. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run [Config Advisor](#).

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Renaming a Brocade FC switch

You might need to rename a Brocade FC switch to ensure consistent naming throughout your configuration.

About this task

[Enable console logging](#) before performing this task.

Steps

1. Persistently disable the switch or switches in one fabric:

```
switchcfgpersistentdisable
```

The following example shows the output for the **switchcfgpersistentdisable** command:

```
7840_FCIP_2:admin> switchcfgpersistentdisable
Switch's persistent state set to 'disabled'
2018/03/09-07:41:06, [ESM-2105], 146080, FID 128, INFO, 7840_FCIP_2, VE
Tunnel 24 is DEGRADED.
2018/03/09-07:41:06, [ESM-2104], 146081, FID 128, INFO, 7840_FCIP_2, VE
Tunnel 24 is OFFLINE.

7840_FCIP_2:admin>
```

2. Rename the switch or switches:

switchname *new-switch-name*

If you are renaming both switches in the fabric, use the same command on each switch.

The following example shows the output for the **switchname *new-switch-name*** command:

```
7840_FCIP_2:admin> switchname FC_switch_1_B
Committing configuration...
Done.
Switch name has been changed.Please re-login into the switch for the
change to be applied.
2018/03/09-07:41:20, [IPAD-1002], 146082, FID 128, INFO, FC_switch_1_B,
Switch name has been successfully changed to FC_switch_1_B.
7840_FCIP_2:admin>
```

3. Reboot the switch or switches:

reboot

If you are renaming both switches in the fabric, reboot both switches. Once the reboot is complete, the switch is renamed in all places.

The following example shows the output for the **reboot** command:

```
7840_FCIP_2:admin> reboot
Warning: This command would cause the switch to reboot
and result in traffic disruption.
Are you sure you want to reboot the switch [y/n]?y
2018/03/09-07:42:08, [RAS-1007], 146083, CHASSIS, INFO, Brocade7840,
System is about to reload.
Rebooting! Fri Mar 9 07:42:11 CET 2018

Broadcast message from root (ttyS0) Fri Mar 9 07:42:11 2018...

The system is going down for reboot NOW !!
INIT: Switching to runlevel: 6
INIT:
2018/03/09-07:50:48, [ESM-1013], 146104, FID 128, INFO, FC_switch_1_B,
DP0 Configuration replay has completed.
2018/03/09-07:50:48, [ESM-1011], 146105, FID 128, INFO, FC_switch_1_B,
DP0 is ONLINE.

*** CORE FILES WARNING (03/09/18 - 08:00:00 ) ***
10248 KBytes in 1 file(s)
use "supportsave" command to upload

*** FFDC FILES WARNING (03/09/18 - 08:00:00 ) ***
520 KBytes in 1 file(s)
```

4. Persistently enable the switches: **switchcfgpersistentenable**

The following example shows the output for the **switchcfgpersistentenable** command:

```

FC_switch_1_B:admin> switchcfgpersistentenable
Switch's persistent state set to 'enabled'
FC_switch_1_B:admin>
FC_switch_1_B:admin>
FC_switch_1_B:admin> 2018/03/09-08:07:07, [ESM-2105], 146106, FID 128,
INFO, FC_switch_1_B, VE Tunnel 24 is DEGRADED.
2018/03/09-08:07:10, [ESM-2106], 146107, FID 128, INFO, FC_switch_1_B,
VE Tunnel 24 is ONLINE.

FC_switch_1_B:admin>

```

```

FC_switch_1_B:admin> switchshow
switchName:      FC_switch_1_B
switchType:      148.0
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     6
switchId:        fffc06
switchWwn:       10:00:50:eb:1a:9a:a5:79
zoning:          ON (CFG_FAB_2_RCF_9_3)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
Address Mode:    0
HIF Mode:        OFF

```

Index	Port	Address	Media	Speed	State	Proto
0	0	060000	id	16G	Online	FC F-Port
		50:0a:09:81:06:a5:5a:08				
1	1	060100	id	16G	Online	FC F-Port
		50:0a:09:83:06:a5:5a:08				

5. Verify that the switch name change is visible from the ONTAP cluster prompt:

storage switch show

The following example shows the output for the **storage switch show** command:

```

cluster_A::*> storage switch show
(storage switch show)
          Symbolic                               Is
Monitor
Switch      Name      Vendor  Model  Switch  WWN          Monitored
Status
-----
-----
Brocade_172.20.7.90
          RTP-FC01-510Q40
          Brocade Brocade7840
          1000c4f57c904bc8 true
ok
Brocade_172.20.7.91
          RTP-FC02-510Q40
          Brocade Brocade7840
          100050eb1a9aa579 true
ok
Brocade_172.20.7.92

```

Disabling encryption on Brocade FC switches

You might need to disable encryption on Brocade FC switches.

Steps

1. Send an AutoSupport message from both sites indicating the beginning of maintenance.

```
cluster_A::> autosupport invoke -node * -type all -message MAINT=4h
```

```
cluster_B::> autosupport invoke -node * -type all -message MAINT=4h
```

2. Verify the operation of the MetroCluster configuration from Cluster A.

- a. Confirm the MetroCluster configuration and that the operational mode is normal:

metrocluster show

```
cluster_A::> metrocluster show
```

- b. Perform a MetroCluster check:

metrocluster check run

```
cluster_A::> metrocluster check run
```

- c. Display the results of the MetroCluster check:

metrocluster check show

```
cluster_A::> metrocluster check show
```

3. Check the status of both switches:

fabric show

```
switch_A_1:admin> fabric show
```

```
switch_B_1:admin> fabric show
```

4. Disable both switches:

switchdisable

```
switch_A_1:admin> switchdisable
```

```
switch_B_1:admin> switchdisable
```

5. Check the available paths for the nodes on each cluster:

sysconfig

```
cluster_A::> system node run -node node-name -command sysconfig -a
```

```
cluster_B::> system node run -node node-name -command sysconfig -a
```

As the switch fabric is now disabled, the System Storage Configuration should be Single-Path HA.

6. Check the aggregate status for both clusters.

```
cluster_A::> aggr status
```

```
cluster_B::> aggr status
```

System output should show the aggregates are mirrored and normal for both clusters:

```
mirrored,normal
```

7. Repeat the following substeps from the admin prompt on both switches.

a. Show which ports are encrypted:

portenccompshow

```
switch_A_1:admin> portenccompshow
```

b. Disable encryption on the encrypted ports:

portcfgencrypt - disable port-number

```
switch_A_1:admin> portcfgencrypt --disable 40
switch_A_1:admin> portcfgencrypt --disable 41
switch_A_1:admin> portcfgencrypt --disable 42
switch_A_1:admin> portcfgencrypt --disable 43
```

c. Set the authentication type to all:

authUtil --set -a all

```
switch_A_1:admin> authUtil --set -a all
```

d. Set the authentication policy on the switch. to off:

authutil --policy -sw off

```
switch_A_1:admin> authutil --policy -sw off
```

e. Set the authentication Diffie-Hellman group to * :

authutil --set -g *

```
switch_A_1:admin> authUtil --set -g *
```

f. Delete the secret key database:

```
secAuthSecret --remove -all
```

```
switch_A_1:admin> secAuthSecret --remove -all
```

g. Confirm that encryption is disabled on the ports:

```
portenccompshow
```

```
switch_A_1:admin> portenccompshow
```

h. Enable the switch:

```
switchenable
```

```
switch_A_1:admin> switchenable
```

i. Confirm the status of the ISLs:

```
islshow
```

```
switch_A_1:admin> islshow
```

8. Check the available paths for the nodes on each cluster:

```
sysconfig
```

```
cluster_A::> system node run -node * -command sysconfig -a
```

```
cluster_B::> system node run -node * -command sysconfig -a
```

The system output should indicate that System Storage Configuration has changed back to Quad-Path HA.

9. Check the aggregate status for both clusters.

```
cluster_A::> aggr status
```

```
cluster_B::> aggr status
```

The system should show that the aggregates are mirrored and normal for both clusters as shown in the following system output:

```
mirrored,normal
```

10. Verify the operation of the MetroCluster configuration from Cluster A.

a. Perform a MetroCluster check:

metrocluster check run

```
cluster_A::> metrocluster check run
```

b. Display the results of the MetroCluster check:

metrocluster check show

```
cluster_A::> metrocluster check show
```

11. Send an AutoSupport message from both sites indicating the end of maintenance.

```
cluster_A::> autosupport invoke -node node-name -type all -message  
MAINT=END
```

```
cluster_B::> autosupport invoke -node node-name -type all -message  
MAINT=END
```

Change ISL properties, ISL ports, or the IOD/OOD configuration on a Brocade switch

You might need to add ISLs to a switch if you are adding or upgrading hardware such as additional or faster controllers or switches.

Before you begin

Ensure that the system is properly configured, that all fabric switches are operational, and that no errors exist.

[Enable console logging](#) before performing this task.

If the equipment on the ISL link changes and the new link configuration no longer supports the current configuration----trunking and ordered delivery----then the fabric needs to be reconfigured for the correct routing policy: either in-order-deliver (IOD) or out-of-order-delivery (OOD).



To make changes to OOD from ONTAP software, use the following steps: [Configuring in-order delivery or out-of-order delivery of frames on ONTAP software](#)

Steps

1. Disable the FCVI and storage HBA ports:

```
portcfgpersistentdisable port number
```

By default the first 8 ports (ports 0 through 7) are used for FCVI and Storage HBA. The ports must be persistently disabled so that the ports remain disabled in the event of a switch reboot.

The following example shows ISL ports 0—7 being disabled on both switches:

```
Switch_A_1:admin> portcfgpersistentdisable 0-7
Switch_B_1:admin> portcfgpersistentdisable 0-7
```

2. Change the ISL ports as required.

Option	Step
To change the speed of an ISL port...	<p>Use the <code>portcfgspeed <i>port number port speed</i></code> command on both switches on the fabric.</p> <p>In the following example, you change the ISL port speed from 40 Gbps to 16 Gbps:</p> <pre>brocade_switch_A_1:admin> portcfgspeed 40 16</pre> <p>You can verify that the speed has changed using the <code>switchshow</code> command:</p> <pre>brocade_switch_A_1:admin> switchshow</pre> <p>You should see the following output:</p> <pre>. . . 40 40 062800 id 16G No_Sync FC Disabled . . .</pre>
To change the distance of an ISL port...	Use the <code>portcfglongdistance <i>port number port distance</i></code> command on both switches in the fabric.
To remove an ISL...	Disconnect the link.
To add an ISL...	Insert SFPs into the ports you are adding as ISL ports. Ensure that these ports are listed in Install a fabric-attached MetroCluster for the switch to which you are adding them.

To relocate an ISL...	Relocating an ISL is the same as removing and then adding an ISL. First, remove the ISL by disconnecting the link and then insert SFPs into the ports you are adding as ISL ports.
-----------------------	--



When you make changes to ISL ports you might also need to apply additional settings recommended by the WDM vendor. Refer to the WDM vendor documentation for guidance.

3. Reconfigure for out-of-order delivery (OOD) or in-order-delivery (IOD).



If the routing policies remain the same, you do not need to reconfigure and this step can be ignored. The ONTAP configuration needs to match the fabric configuration. If the fabric is configured for OOD, then ONTAP must also be configured for OOD. The same applies for IOD.

This step should be executed in the following scenarios:

- More than one ISL formed a trunk before the change, but after the change, trunking is no longer supported. In this case, you must configure the fabric for OOD.
- There is one ISL before the change and multiple ISLs after the change.
- If multiple ISLs form a trunk, configure the fabric for IOD.
If multiple ISLs **cannot** form a trunk, configure the fabric for OOD.
- Persistently disable the switches using the `switchcfgpersistentdisable` command as shown in the following example:

```
Switch_A_1:admin> switchcfgpersistentdisable
Switch_B_1:admin> switchcfgpersistentdisable
```

- a. Configure the trunking mode for each ISL `portcfgtrunkport port number` as shown in the following table:

Scenario	Steps
Configure the ISL for trunking \(\IOD\)	<p>Set the <code>portcfgtrunkport port number</code> to 1:</p> <pre>FC_switch_A_1:admin> portcfgtrunkport 20 1 FC_switch_A_1:admin> portcfgtrunkport 21 1 FC_switch_B_1:admin> portcfgtrunkport 20 1 FC_switch_B_1:admin> portcfgtrunkport 21 1</pre>

Configure the ISL for trunking \((OOD\)	<p>Set the portcfgtrunkport <i>port number</i> to 0:</p> <pre> FC_switch_A_1:admin> portcfgtrunkport 20 0 FC_switch_A_1:admin> portcfgtrunkport 21 0 FC_switch_B_1:admin> portcfgtrunkport 20 0 FC_switch_B_1:admin> portcfgtrunkport 21 0 </pre>
---	---

b. Configure the fabric for IOD or OOD as required.

Scenario	Steps
Configure the fabric for IOD	<p>Set the three settings of IOD, APT, and DLS using the <code>iodset</code>, <code>aptpolicy</code>, and <code>dlsreset</code> commands as shown in the following example:</p> <pre> Switch_A_1:admin> iodset Switch_A_1:admin> aptpolicy 1 Policy updated successfully. Switch_A_1:admin> dlsreset FC_switch_A_1:admin>portcfgtrunkport 40 1 FC_switch_A_1:admin>portcfgtrunkport 41 1 Switch_B_1:admin> iodset Switch_B_1:admin> aptpolicy 1 Policy updated successfully. Switch_B_1:admin> dlsreset FC_switch_B_1:admin>portcfgtrunkport 20 1 FC_switch_B_1:admin>portcfgtrunkport 21 1 </pre>

Configure the fabric for OOD

Set the three settings of IOD, APT, and DLS using the `iodreset`, `aptpolicy`, and `dlset` commands as shown in the following example:

```
Switch_A_1:admin> iodreset
Switch_A_1:admin> aptpolicy 3
Policy updated successfully.
Switch_A_1:admin> dlset
FC_switch_A_1:admin> portcfgtrunkport 40 0
FC_switch_A_1:admin> portcfgtrunkport 41 0

Switch_B_1:admin> iodreset
Switch_B_1:admin> aptpolicy 3
Policy updated successfully.
Switch_B_1:admin> dlset
FC_switch_B_1:admin> portcfgtrunkport 40 0
FC_switch_B_1:admin> portcfgtrunkport 41 0
```

c. Enable the switches persistently:

```
switchcfgpersistentenable
```

```
switch_A_1:admin>switchcfgpersistentenable
switch_B_1:admin>switchcfgpersistentenable
```

If this command does not exist, use the `switchenable` command as shown in the following example:

```
brocade_switch_A_1:admin>
switchenable
```

d. Verify the OOD settings using the `iodshow`, `aptpolicy`, and `dlshow` commands as shown in the following example:

```
switch_A_1:admin> iodshow
IOD is not set

switch_A_1:admin> aptpolicy

Current Policy: 3 0(ap)

3 0(ap) : Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
0: AP Shared Link Policy
1: AP Dedicated Link Policy
command aptpolicy completed

switch_A_1:admin> dlsshow
DLS is set by default with current routing policy
```



You must run these commands on both switches.

- e. Verify the IOD settings using the `iodshow`, `aptpolicy`, and `dlsshow` commands as shown in the following example:

```
switch_A_1:admin> iodshow
IOD is set

switch_A_1:admin> aptpolicy
Current Policy: 1 0(ap)

3 0(ap) : Default Policy
1: Port Based Routing Policy
3: Exchange Based Routing Policy
0: AP Shared Link Policy
1: AP Dedicated Link Policy
command aptpolicy completed

switch_A_1:admin> dlsshow
DLS is not set
```



You must run these commands on both switches.

4. Verify that the ISLs are online and trunked (if the linking equipment supports trunking) using the `islshow` and `trunkshow` commands.



If FEC is enabled, the deskew value of the last online port of the trunk group might show a difference of up to 36 although the cables are all of the same length.

Are ISLs trunked?	You see the following system output...
Yes	<p>If the ISLs are trunked, only a single ISL appears in the output for the <code>islshow</code> command. Either port 40 or 41 can appear depending on which is the trunk master. The output of <code>trunkshow</code> should one trunk with ID "1" listing both the physical ISLs on ports 40 and 41. In the following example the ports 40 and 41 are configured for use as an ISL:</p> <pre data-bbox="451 489 1481 785">switch_A_1:admin> islshow 1: 40-> 40 10:00:00:05:33:88:9c:68 2 switch_B_1 sp: 16.000G bw: 32.000G TRUNK CR_RECOV FEC switch_A_1:admin> trunkshow 1: 40-> 40 10:00:00:05:33:88:9c:68 2 deskew 51 MASTER 41-> 41 10:00:00:05:33:88:9c:68 2 deskew 15</pre>
No	<p>If the ISLs are not trunked, both ISLs appear separately in the outputs for <code>islshow</code> and <code>trunkshow</code>. Both commands list the ISLs with their ID of "1" and "2". In the following example, the ports "40" and "41" are configured for use as an ISL:</p> <pre data-bbox="451 978 1481 1352">switch_A_1:admin> islshow 1: 40-> 40 10:00:00:05:33:88:9c:68 2 switch_B_1 sp: 16.000G bw: 16.000G TRUNK CR_RECOV FEC 2: 41-> 41 10:00:00:05:33:88:9c:68 2 switch_B_1 sp: 16.000G bw: 16.000G TRUNK CR_RECOV FEC switch_A_1:admin> trunkshow 1: 40-> 40 10:00:00:05:33:88:9c:68 2 deskew 51 MASTER 2: 41-> 41 10:00:00:05:33:88:9c:68 2 deskew 48 MASTER</pre>

5. Run the `spinfab` command on both switches to verify that the ISLs are healthy:

```
switch_A_1:admin> spinfab -ports 0/40 - 0/41
```

6. Enable the ports that were disabled in step 1:

```
portenable port number
```

The following example shows ISL ports "0" through "7" being enabled:

```
brocade_switch_A_1:admin> portenable 0-7
```

Replacing a Cisco FC switch

You must use Cisco-specific steps to replace a failed Cisco FC switch.

Before you begin

You need the admin password and access to an FTP or SCP server.

[Enable console logging](#) before performing this task.

About this task

This procedure is nondisruptive and takes approximately two hours to complete.

In the examples in this procedure, FC_switch_A_1 is the healthy switch and FC_switch_B_1 is the impaired switch. The switch port usage in the examples is shown in the following table:

Role	Ports
FC-VI connections	1, 4
HBA connections	2, 3, 5, 6
FC-to-SAS bridge connections	7, 8
ISL connections	36, 40

The examples show two FC-to-SAS bridges. If you have more, you must disable and subsequently enable the additional ports.

Your switch port usage should follow the recommended assignments.

- [Port assignments for FC switches](#)

Steps

1. Disable the ISL ports on the healthy switch to fence off the impaired switch.

These steps are performed on the healthy switch.

- a. Enter configuration mode:

```
conf t
```

- b. Disable the ISL ports on the healthy switch with the `interface` and `shut` commands.

```
FC_switch_A_1# conf t
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/40
FC_switch_A_1(config)# shut
```

- c. Exit configuration mode and copy the configuration to the startup configuration.

```
FC_switch_A_1(config)# end
FC_switch_A_1# copy running-config startup-config
FC_switch_A_1#
```

2. Fence off the FC-VI and HBA ports on the impaired switch (if it is still running).

These steps are performed on the impaired switch.

a. Enter configuration mode:

```
conf t
```

b. If the impaired switch is still operational, disable the FC-VI and HBA ports on the impaired switch with the interface and shut commands.

```
FC_switch_B_1(config)# interface fc1/1
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/4
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/2-3
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/5-6
FC_switch_B_1(config)# shut
```

c. Exit configuration mode and copy the configuration to the startup configuration.

```
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
FC_switch_B_1#
```

3. If the impaired switch is still operational, determine the WWN for the switch:

```
show wwn switch
```

```
FC_switch_B_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:e3:86:50
FC_switch_B_1#
```

4. Boot and preconfigure the replacement switch, prior to physically installing it.

At this point the replacement switch is not cabled to the MetroCluster configuration. The ISL ports on the partner switch are disabled (in shut mode) and offline.

a. Power on the replacement switch and let it boot up.

- b. Check the firmware version on the replacement switch to confirm that it matches the version of the other FC switches:

```
show version
```

- c. Configure the replacement switch as described in the *MetroCluster Installation and Configuration Guide*, skipping the “Configuring zoning on a Cisco FC switch” section.

Fabric-attached MetroCluster installation and configuration

You will configure zoning later in this procedure.

- d. Disable the FC-VI, HBA, and storage ports on the replacement switch.

```
FC_switch_B_1# conf t
FC_switch_B_1(config)# interface fc1/1
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/4
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/2-3
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/5-6
FC_switch_B_1(config)# shut
FC_switch_B_1(config)# interface fc1/7-8
FC_switch_B_1(config)# shut
FC_switch_B_1# copy running-config startup-config
FC_switch_B_1#
```

5. Physically replace the impaired switch:
 - a. Power off the impaired switch.
 - b. Power off the replacement switch.
 - c. Uncable and remove the impaired switch, carefully noting which cables connected to which ports.
 - d. Install the replacement switch in the rack.
 - e. Cable the replacement switch exactly as the impaired switch was cabled.
 - f. Power on the replacement switch.
6. Enable the ISL ports on the replacement switch.

```
FC_switch_B_1# conf t
FC_switch_B_1(config)# interface fc1/36
FC_switch_B_1(config)# no shut
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
FC_switch_B_1(config)# interface fc1/40
FC_switch_B_1(config)# no shut
FC_switch_B_1(config)# end
FC_switch_B_1#
```

7. Verify that the ISL ports on the replacement switch are up:

```
show interface brief
```

8. Adjust the zoning on the replacement switch to match the MetroCluster configuration:

a. Distribute the zoning information from the healthy fabric.

In this example, FC_switch_B_1 has been replaced and the zoning information is retrieved from FC_switch_A_1:

```
FC_switch_A_1(config-zone)# zoneset distribute full vsan 10
FC_switch_A_1(config-zone)# zoneset distribute full vsan 20
FC_switch_A_1(config-zone)# end
```

b. On the replacement switch, verify that the zoning information was properly retrieved from the healthy switch:

```
show zone
```

```

FC_switch_B_1# show zone
zone name FC-VI_Zone_1_10 vsan 10
  interface fc1/1 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/4 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/1 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/4 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25A vsan 20
  interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/3 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/6 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/3 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/6 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25B vsan 20
  interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/3 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/6 swwn 20:00:54:7f:ee:e3:86:50
  interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/3 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/6 swwn 20:00:54:7f:ee:b8:24:c0
FC_switch_B_1#

```

c. Find the WWNs of the switches.

In this example, the two switch WWNs are as follows:

- FC_switch_A_1: 20:00:54:7f:ee:b8:24:c0
- FC_switch_B_1: 20:00:54:7f:ee:c6:80:78

```

FC_switch_B_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:c6:80:78
FC_switch_B_1#

FC_switch_A_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:b8:24:c0
FC_switch_A_1#

```

a. Remove zone members that do not belong to the switch WWNs of the two switches.

In this example, “no member interface” in the output shows that the following members are not associated with the switch WWN of either of the switches in the fabric and must be removed:

- zone name FC-VI_Zone_1_10 vsan 10
 - interface fc1/1 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
 - zone name STOR_Zone_1_20_25A vsan 20
 - interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
 - zone name STOR_Zone_1_20_25B vsan 20
 - interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
 - interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
- The following example shows the removal of these interfaces:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# zone name FC-VI_Zone_1_10 vsan 10
FC_switch_B_1(config-zone)# no member interface fc1/1 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/2 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25A vsan
20
FC_switch_B_1(config-zone)# no member interface fc1/5 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/8 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/9 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/10 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/11 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25B vsan
20
FC_switch_B_1(config-zone)# no member interface fc1/8 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/9 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/10 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/11 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# save running-config startup-config
FC_switch_B_1(config-zone)# zoneset distribute full 10
FC_switch_B_1(config-zone)# zoneset distribute full 20
FC_switch_B_1(config-zone)# end
FC_switch_B_1# copy running-config startup-config

```

b. Add the ports of the replacement switch to the zones.

All the cabling on the replacement switch must be the same as on the impaired switch:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# zone name FC-VI_Zone_1_10 vsan 10
FC_switch_B_1(config-zone)# member interface fc1/1 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/2 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25A vsan 20
FC_switch_B_1(config-zone)# member interface fc1/5 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/8 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/9 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/10 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/11 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25B vsan 20
FC_switch_B_1(config-zone)# member interface fc1/8 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/9 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/10 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/11 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# save running-config startup-config
FC_switch_B_1(config-zone)# zoneset distribute full 10
FC_switch_B_1(config-zone)# zoneset distribute full 20
FC_switch_B_1(config-zone)# end
FC_switch_B_1# copy running-config startup-config

```

c. Verify that the zoning is properly configured:

```
show zone
```

The following example output shows the three zones:

```

FC_switch_B_1# show zone
zone name FC-VI_Zone_1_10 vsan 10
  interface fc1/1 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/2 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/1 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25A vsan 20
  interface fc1/5 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/8 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/9 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/10 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/11 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25B vsan 20
  interface fc1/8 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/9 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/10 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/11 swwn 20:00:54:7f:ee:c6:80:78
  interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
  interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0
FC_switch_B_1#

```

- d. Enable the connectivity to storage and the controllers.

The following example shows the port usage:

```

FC_switch_A_1# conf t
FC_switch_A_1(config)# interface fc1/1
FC_switch_A_1(config)# no shut
FC_switch_A_1(config)# interface fc1/4
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/2-3
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/5-6
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# interface fc1/7-8
FC_switch_A_1(config)# shut
FC_switch_A_1# copy running-config startup-config
FC_switch_A_1#

```

9. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Change ISL properties and the IOD/OOD configuration on a Cisco FC switch

You can add Inter-Switch Links (ISLs), change ISL speed, and reconfigure in-order delivery (IOD) or out of-order delivery (OOD) settings on a Cisco FC switch.

Add ISLs to a Cisco FC switch

You might need add ISLs to a switch if you are adding or upgrading hardware, for example, adding or upgrading to faster controllers or faster switches.

About this task

Perform these steps on both switches in the fabric to verify ISL connectivity.

Steps

1. Disable the ISL ports of the ISLs to be added on both switches in the fabric:

```
FC_switch_A_1#config t
```

Enter the following configuration commands, one per line. Enter CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# shut
FC_switch_A_1(config)# end
```

2. Insert SFPs into the ports you are adding as ISL ports, and cable them according to [Cable a fabric-attached MetroCluster configuration](#).

Verify that these ports are listed in the cabling documentation for the switch model that you are adding them to.

3. Configure the ISL ports by following the steps in [Cabling the ISLs between MetroCluster sites](#).
4. Enable all ISL ports (if not enabled) on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1# interface fc1/36
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config)# end
```

5. Verify that the ISLs are established between both switches:

```
show topology isl
```

6. Repeat the procedure on the second fabric:

```

-----
-----

```

	Local				Remote				VSAN	Cost	I/F	PC
I/F	Band	PC Domain	SwName	Port	Port	SwName	Domain	PC			Stat	Stat
	Speed	width										
	1	0x11	cisco9	fc1/36	fc1/36	cisco9	0xbc	1	1	15	up	up
16g	64g											
	1	0x11	cisco9	fc1/40	fc1/40	cisco9	0xbc	1	1	15	up	up
16g	64g											
	1	0x11	cisco9	fc1/44	fc1/44	cisco9	0xbc	1	1	15	up	up
16g	64g											
	1	0x11	cisco9	fc1/48	fc1/48	cisco9	0xbc	1	1	15	up	up
16g	64g											

```

-----
-----

```

Change ISL port speeds on a Cisco FC switch

You can change the speed of ISL ports on a switch to improve the quality of the ISL, for example, lowering the speed on ISLs traveling a greater distance.

About this task

Perform these steps on both switches in the fabric to verify ISL connectivity.

Steps

1. Disable the ISL ports for the ISLs that you want to change the speed for on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# shut
FC_switch_A_1(config)# end
```

2. Change the speed of the ISL ports on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# switchport speed 16000
```



Speeds for the ports are 16 = 16,000 Gbps, 8 = 8,000 Gbps, and 4 = 4,000 Gbps.

Verify that the ISL ports for your switch are listed in [Install a fabric-attached MetroCluster configuration](#).

3. Enable all ISL ports (if not enabled) on both switches in the fabric:

```
FC_switch_A_1# config t
```

Enter the following configuration commands, one per line. End with CTRL-Z after you have entered all of the configuration commands.

```
FC_switch_A_1(config)# interface fc1/36
FC_switch_A_1(config-if)# no shut
FC_switch_A_1(config)# end
```

4. Verify that the ISLs are established between both switches:

```
show topology isl
```

```
-----
-----
          _____Local_____          _____Remote_____  VSAN Cost I/F  PC
I/F  Band
      PC Domain SwName   Port   Port   SwName Domain PC           Stat Stat
Speed width
-----
-----
      1   0x11 cisco9 fc1/36  fc1/36 cisco9 0xbc    1    1   15 up   up
16g   64g
      1   0x11 cisco9 fc1/40  fc1/40 cisco9 0xbc    1    1   15 up   up
16g   64g
      1   0x11 cisco9 fc1/44  fc1/44 cisco9 0xbc    1    1   15 up   up
16g   64g
      1   0x11 cisco9 fc1/48  fc1/48 cisco9 0xbc    1    1   15 up   up
16g   64g
```

5. Repeat the procedure for the second switch fabric.

Reconfigure the VSAN to guarantee IOD or OOD of frames

The standard IOD settings are recommended. You should only reconfigure OOD if necessary.

Reconfigure IOD

Perform the following step to reconfigure IOD of frames.

Steps

1. Enter configuration mode:

```
conf t
```

2. Enable the in-order guarantee of exchanges for the VSAN:

```
in-order-guarantee vsan <vsan-ID>
```



For FC-VI VSANs (FCVI_1_10 and FCVI_2_30), you must enable in-order guarantee of frames and exchanges only on VSAN 10.

- a. Enable load balancing for the VSAN:

```
vsan <vsan-ID> loadbalancing src-dst-id
```

- b. Exit configuration mode:

```
end
```

- c. Copy the running-config to the startup-config:

```
copy running-config startup-config
```

The commands to configure IOD of frames on FC_switch_A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# in-order-guarantee vsan 10
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to configure IOD of frames on FC_switch_B_1:

```
FC_switch_B_1# config t
FC_switch_B_1(config)# in-order-guarantee vsan 10
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config
```

Reconfigure OOD

Perform the following steps to reconfigure OOD of frames.

Steps

1. Enter configuration mode:

```
conf t
```

2. Disable the in-order guarantee of exchanges for the VSAN:

```
no in-order-guarantee vsan <vsan-ID>
```

3. Enable load balancing for the VSAN:

```
vsan <vsan-ID> loadbalancing src-dst-id
```

4. Exit configuration mode:

```
end
```

5. Copy the running-config to the startup-config:

```
copy running-config startup-config
```

The commands to configure OOD of frames on FC_switch_A_1:

```
FC_switch_A_1# config t
FC_switch_A_1(config)# no in-order-guarantee vsan 10
FC_switch_A_1(config)# vsan database
FC_switch_A_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_A_1(config-vsan-db)# end
FC_switch_A_1# copy running-config startup-config
```

The commands to configure OOD of frames on FC_switch_B_1:

```
FC_switch_B_1# config t
FC_switch_B_1(config)# no in-order-guarantee vsan 10
FC_switch_B_1(config)# vsan database
FC_switch_B_1(config-vsan-db)# vsan 10 loadbalancing src-dst-id
FC_switch_B_1(config-vsan-db)# end
FC_switch_B_1# copy running-config startup-config
```



When configuring ONTAP on the controller modules, OOD must be explicitly configured on each controller module in the MetroCluster configuration.

[Learn about configuring IOD or OOD of frames on ONTAP software.](#)

Change the vendor or model of FC switches

You might need to change the vendor for FC switches from Cisco to Brocade or vice versa, change the switch model, or change both.

About this task

- This procedure applies when you are using NetApp validated switches.
- [Enable console logging](#) before performing this task.
- You must perform the steps in this procedure on one Fabric at a time, for both Fabrics in the configuration.

Steps

1. Check the health of the configuration.
 - a. Check that the MetroCluster is configured and in normal mode on each cluster: **metrocluster show**

```
cluster_A::> metrocluster show
Cluster                               Entry Name                               State
-----                               -
Local: cluster_A                       Configuration state configured
Mode                                    normal
AUSO Failure Domain auso-on-cluster-
disaster
Remote: cluster_B                       Configuration state configured
Mode                                    normal
AUSO Failure Domain auso-on-cluster-
disaster
```

- b. Check that mirroring is enabled on each node: **metrocluster node show**

```
cluster_A::> metrocluster node show
DR                                     Configuration  DR
Group Cluster Node                     State          Mirroring Mode
-----
1      cluster_A
      node_A_1      configured    enabled    normal
      cluster_B
      node_B_1      configured    enabled    normal
2 entries were displayed.
```

- c. Check that the MetroCluster components are healthy: **metrocluster check run**

```
cluster_A::> metrocluster check run
```

```
Component          Result
-----
nodes              ok
lifs               ok
config-replication ok
aggregates        ok
4 entries were displayed.
```

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results.

To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

- d. Check that there are no health alerts: **system health alert show**
2. Configure the new switches before installation.

Follow the steps in [Configure the FC switches](#).
3. Disconnect the connections from the old switches by removing the connections in the following order:
 - a. Disconnect the MetroCluster FC and FCVI interfaces.
 - b. Disconnect the ATTO FibreBridge bridges.
 - c. Disconnect the MetroCluster ISLs.
4. Power off the old switches, remove the cables, and physically replace the old switches with the new switch.
5. Cable the switches in the following order:

You must follow the steps in [Cabling a fabric-attached MetroCluster configuration](#).

- a. Cable the ISLs to the remote site.
 - b. Cable the ATTO FibreBridge bridges.
 - c. Cable the MetroCluster FC and FCVI interfaces.
6. Power up the switches.
 7. Verify that the MetroCluster configuration is healthy by repeating [Step 1](#).
 8. Repeat Step 1 to Step 7 for the second Fabric in the configuration.

Replacing a shelf nondisruptively in a fabric-attached MetroCluster configuration

You might need to know how to replace a shelf nondisruptively in a fabric-attached MetroCluster configuration.



This procedure is only for use in a fabric-attached MetroCluster configuration.

Disabling access to the shelf

You must disable access to the shelf before you replace the shelf modules.

Check the overall health of the configuration. If the system does not appear healthy, address the issue first before proceeding.

Steps

1. From both clusters, offline all plexes with disks on the affected shelf stack:

```
aggr offline plex_name
```

The example shows the commands for offlining plexes for a controller running ONTAP.

```
cluster_A_1::> storage aggregate plex offline -aggr aggrA_1_0 -plex
plex0
cluster_A_1::> storage aggregate plex offline -aggr dataA_1_data -plex
plex0
cluster_A_2::> storage aggregate plex offline -aggr aggrA_2_0 -plex
plex0
cluster_A_2::> storage aggregate plex offline -aggr dataA_2_data -plex
plex0
```

2. Verify that the plexes are offline:

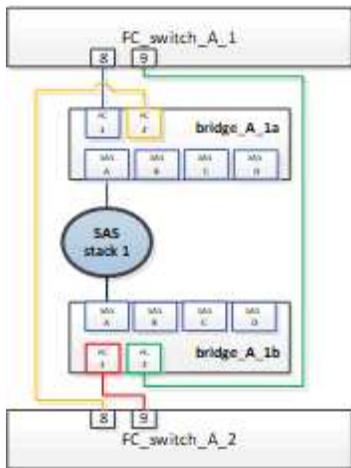
```
aggr status -raggr_name
```

The example shows the commands for verifying that the aggregates are offline for a controller running cMode.

```
Cluster_A_1::> storage aggregate show -aggr aggrA_1_0
Cluster_A_1::> storage aggregate show -aggr dataA_1_data
Cluster_A_2::> storage aggregate show -aggr aggrA_2_0
Cluster_A_2::> storage aggregate show -aggr dataA_2_data
```

3. Disable the SAS ports or switch ports depending on whether the bridges connecting the target shelf are connecting a single SAS stack or two or more SAS stacks:
 - If the bridges are connecting a single SAS stack, disable the switch ports that the bridges are connected to using the appropriate command for your switch.

The following example shows a pair of bridges that connect a single SAS stack, which contains the target shelf:



Switch ports 8 and 9 on each switch connect the bridges to the network.

The following example shows ports 8 and 9 being disabled on a Brocade switch.

```
FC_switch_A_1:admin> portDisable 8
FC_switch_A_1:admin> portDisable 9

FC_switch_A_2:admin> portDisable 8
FC_switch_A_2:admin> portDisable 9
```

The following example shows port 8 and 9 being disabled on a Cisco switch.

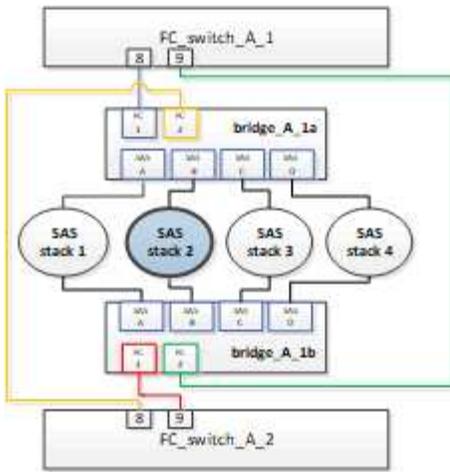
```
FC_switch_A_1# conf t
FC_switch_A_1(config)# int fc1/8
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# int fc1/9
FC_switch_A_1(config)# shut
FC_switch_A_1(config)# end

FC_switch_A_2# conf t
FC_switch_A_2(config)# int fc1/8
FC_switch_A_2(config)# shut
FC_switch_A_2(config)# int fc1/9
FC_switch_A_2(config)# shut
FC_switch_A_2(config)# end
```

- If the bridges are connecting two or more SAS stacks, disable the SAS ports connecting the bridges to the target shelf:

`SASportDisable port number`

The following example shows a pair of bridges that connect four SAS stacks. SAS stack 2 contains the target shelf:



SAS port B connects the bridges to the target shelf. By disabling only SAS port B on both shelves, the other SAS stacks can continue to serve data during the replacement procedure.

In this case, disable the SAS port connecting the bridge to the target shelf:

```
SASportDisable port number
```

The following example shows SAS port B being disabled from the bridge and also verifies that it is disabled. You must repeat the command on both bridges.

```
Ready. *
SASPortDisable B

SAS Port B has been disabled.
```

4. If you previously disabled the switch ports, verify that they are disabled:

```
switchShow
```

The example shows that the switch ports are disabled on a Brocade switch.

```
FC_switch_A_1:admin> switchShow
FC_switch_A_2:admin> switchShow
```

The example shows that the switch ports are disabled on a Cisco switch.

```
FC_switch_A_1# show interface fc1/6
FC_switch_A_2# show interface fc1/6
```

5. Wait for ONTAP to realize that the disk is missing.
6. Power off the shelf that you want to replace.

Replacing the shelf

You must physically remove all of the cables and the shelf before inserting and cabling the new shelf and shelf modules.

Steps

1. Remove all disks and disconnect all cables from the shelf that is being replaced.
2. Remove the shelf modules.
3. Insert the new shelf.
4. Insert the new disks into the new shelf.
5. Insert the shelf modules.
6. Cable the shelf (SAS or Power).
7. Power on the shelf.

Reenabling access and verifying the operation

After the shelf has been replaced, you need to reenable access and verify that the new shelf is operating correctly.

Steps

1. Verify that the shelf powers properly and the links on the IOM modules are present.
2. Enable the switch ports or SAS port according to the following scenarios:

Option	Step
--------	------

If you previously disabled switch ports

- a. Enable the switch ports:

```
portEnable port number
```

The example shows the switch port being enabled on a Brocade switch.

```
Switch_A_1:admin> portEnable 6  
Switch_A_2:admin> portEnable 6
```

The example shows the switch port being enabled on a Cisco switch.

```
Switch_A_1# conf t  
Switch_A_1(config)# int fc1/6  
Switch_A_1(config)# no shut  
Switch_A_1(config)# end  
  
Switch_A_2# conf t  
Switch_A_2(config)# int fc1/6  
Switch_A_2(config)# no shut  
Switch_A_2(config)# end
```

If you previously disabled a SAS port

- a. Enable the SAS port connecting the stack to the shelf location:

```
SASportEnable port number
```

The example shows SAS port A being enabled from the bridge and also verifies that it is enabled.

```
Ready. *  
SASPortEnable A  
  
SAS Port A has been enabled.
```

3. If you previously disabled the switch ports, verify that they are enabled and online and that all devices are logged in correctly:

```
switchShow
```

The example shows the `switchShow` command for verifying that a Brocade switch is online.

```
Switch_A_1:admin> SwitchShow  
Switch_A_2:admin> SwitchShow
```

The example shows the `switchShow` command for verifying that a Cisco switch is online.

```
Switch_A_1# show interface fc1/6
Switch_A_2# show interface fc1/6
```



After several minutes, ONTAP detects that new disks have been inserted and displays a message for each new disk.

4. Verify that the disks have been detected by ONTAP:

```
sysconfig -a
```

5. Online the plexes that were offline earlier:

```
aggr onlineplex_name
```

The example shows the commands for placing plexes on a controller running cMode back online.

```
Cluster_A_1::> storage aggregate plex online -aggr aggr1 -plex plex2
Cluster_A_1::> storage aggregate plex online -aggr aggr2 -plex plex6
Cluster_A_1::> storage aggregate plex online -aggr aggr3 -plex plex1
```

The plexes begin to resynchronize.



You can monitor the progress of resynchronization using the `aggr status -raggr_name` command.

Hot add storage to a MetroCluster FC configuration

Hot-adding a SAS disk shelf in a direct-attached MetroCluster FC configuration using SAS optical cables

You can use SAS optical cables to hot-add a SAS disk shelf to an existing stack of SAS disk shelves in a direct-attached MetroCluster FC configuration, or as a new stack to a SAS HBA or an onboard SAS port on the controller.

- This procedure is nondisruptive and takes approximately two hours to complete.
- You need the admin password and access to an FTP or SCP server.

This task applies to a MetroCluster FC configuration in which the storage is connected directly to the storage controllers with SAS cables. It does not apply to MetroCluster FC configurations using FC-to-SAS bridges or FC switch fabrics.

Steps

1. Follow the instructions for hot-adding a SAS disk shelf in the *Installation Guide* for your disk shelf model to perform the following tasks to hot-add a disk shelf:

- a. Install a disk shelf for a hot-add.
- b. Turn on the power supplies and set the shelf ID for a hot-add.
- c. Cable the hot-added disk shelf.
- d. Verify SAS connectivity.

Hot add SAS storage to a bridge-attached MetroCluster FC configuration

Hot-adding a stack of SAS disk shelves to an existing pair of FibreBridge 7600N or 7500N bridges

You can hot-add a stack of SAS disk shelves to an existing pair of FibreBridge 7600N or 7500N bridges that have available ports.

Before you begin

- You must have downloaded the latest disk and disk shelf firmware.
- All of the disk shelves in the MetroCluster configuration (existing shelves) must be running the same firmware version. If one or more of the disks or shelves are not running the latest firmware version, update the firmware before attaching the new disks or shelves.

[NetApp Downloads: Disk Drive Firmware](#)

[NetApp Downloads: Disk Shelf Firmware](#)

- The FibreBridge 7600N or 7500N bridges must be connected and have available SAS ports.

About this task

This procedure is written with the assumption that you are using the recommended bridge management interfaces: the ATTO ExpressNAV GUI and the ATTO QuickNAV utility.

You can use the ATTO ExpressNAV GUI to configure and manage a bridge, and to update the bridge firmware. You can use the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port.

You can use other management interfaces, if required. These options include using a serial port or Telnet to configure and manage a bridge and to configure the Ethernet management 1 port, and using FTP to update the bridge firmware. If you choose any of these management interfaces, you must meet the applicable requirements in [Other bridge management interfaces](#).



If you insert a SAS cable into the wrong port, when you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that the cable has been moved to another port.

Steps

1. Properly ground yourself.
2. From the console of either controller, verify that your system has disk autoassignment enabled:

```
storage disk option show
```

The Auto Assign column indicates whether disk autoassignment is enabled.

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
node_A_1	on	on	on	default
node_A_2	on	on	on	default
2 entries were displayed.				

3. On each bridge in the pair, enable the SAS port that will connect to the new stack:

```
SASPortEnable port-letter
```

The same SAS port (B, C, or D) must be used on both bridges.

4. Save the configuration and reboot each bridge:

```
SaveConfiguration Restart
```

5. Cable the disk shelves to the bridges:

- a. Daisy-chain the disk shelves in each stack.

The *Installation and Service Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

- b. For each stack of disk shelves, cable IOM A of the first shelf to SAS port A on FibreBridge A, and then cable IOM B of the last shelf to SAS port A on FibreBridge B

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

Each bridge has one path to its stack of disk shelves; bridge A connects to the A-side of the stack through the first shelf, and bridge B connects to the B-side of the stack through the last shelf.



The bridge SAS port B is disabled.

6. Verify that each bridge can detect all of the disk drives and disk shelves to which the bridge is connected.

If you are using the...	Then...
-------------------------	---------

ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of a bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge home page, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page appears with a menu to the left.</p> <p>c. Click Advanced in the menu.</p> <p>d. View the connected devices:</p> <pre>sastargets</pre> <p>e. Click Submit.</p>
Serial port connection	<p>View the connected devices:</p> <pre>sastargets</pre>

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.



If the text “response truncated” appears at the beginning of the output, you can use Telnet to connect to the bridge, and then view all of the output by using the `sastargets` command.

The following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

7. Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
----------------------------	----------------

Correct	Repeat the previous step for each remaining bridge.
Not correct	<ol style="list-style-type: none"> a. Check for loose SAS cables or correct the SAS cabling by repeating the step to cable the disk shelves to the bridges. b. Repeat the previous step for each remaining bridge.

8. Update the disk drive firmware to the most current version from the system console:

```
disk_fw_update
```

You must run this command on both controllers.

[NetApp Downloads: Disk Drive Firmware](#)

9. Update the disk shelf firmware to the most current version by using the instructions for the downloaded firmware.

You can run the commands in the procedure from the system console of either controller.

[NetApp Downloads: Disk Shelf Firmware](#)

10. If your system does not have disk autoassignment enabled, assign disk drive ownership.

[Disk and aggregate management](#)



If you are splitting the ownership of a single stack of disk shelves among multiple controllers, you must disable disk autoassignment (`storage disk option modify -autoassign off *` from both nodes in the cluster) before assigning disk ownership; otherwise, when you assign any single disk drive, the remaining disk drives might be automatically assigned to the same controller and pool.



You must not add disk drives to aggregates or volumes until after the disk drive firmware and disk shelf firmware have been updated and the verification steps in this task have been completed.

11. Verify the operation of the MetroCluster configuration in ONTAP:

a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

b. Check for any health alerts on both clusters:

```
system health alert show
```

c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

d. Perform a MetroCluster check:

```
metrocluster check run
```

e. Display the results of the MetroCluster check:

```
metrocluster check show
```

f. Check for any health alerts on the bridges after adding the new stacks:

```
storage bridge show
```

g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

12. If applicable, repeat this procedure for the partner site.

Hot-adding a stack of SAS disk shelves and bridges to a MetroCluster system

You can hot-add (nondisruptively add) an entire stack, including the bridges, to the MetroCluster system. There must be available ports on the FC switches and you must update switch zoning to reflect the changes.

About this task

- This procedure can be used to add a stack using FibreBridge 7600N or 7500N bridges.
- This procedure is written with the assumption that you are using the recommended bridge management interfaces: the ATTO ExpressNAV GUI and the ATTO QuickNAV utility.
 - You use the ATTO ExpressNAV GUI to configure and manage a bridge, and to update the bridge firmware. You use the ATTO QuickNAV utility to configure the bridge Ethernet management 1 port.
 - You can use other management interfaces, if needed. These options include using a serial port or Telnet to configure and manage a bridge, and to configure the Ethernet management 1 port, and using FTP to update the bridge firmware. If you choose any of these management interfaces, your system must meet the applicable requirements in [Other bridge management interfaces](#)

Preparing to hot-add a stack of SAS disk shelves and bridges

Preparing to hot-add a stack of SAS disk shelves and a pair of bridges involves downloading documents as well as the disk drive and disk shelf firmware.

Before you begin

- Your system must be a supported configuration and must be running a supported version of ONTAP.

[NetApp Interoperability Matrix Tool](#)

- All disk drives and disk shelves in the system must be running the latest firmware version.

You might want to update the disk and shelf firmware throughout the MetroCluster configuration prior to adding shelves.

[Upgrade, revert, or downgrade](#)

- Each FC switch must have one FC port available for one bridge to connect to it.



You might need to upgrade the FC switch depending on the FC switch compatibility.

- The computer you are using to set up the bridges must be running an ATTO supported web browser to use the ATTO ExpressNAV GUI: Internet Explorer 8 or 9, or Mozilla Firefox 3.

The *ATTO Product Release Notes* have an up-to-date list of supported web browsers. You can access this document using the information in the steps.

Steps

1. Download or view the following documents from the NetApp Support Site:
 - [NetApp Interoperability Matrix Tool](#)
 - The *Installation and Service Guide* for your disk shelf model.
2. Download content from the ATTO website and from the NetApp website:
 - a. Go to the ATTO FibreBridge Description page.
 - b. Using the link on the ATTO FibreBridge Description page, access the ATTO web site and download the following:
 - *ATTO FibreBridge Installation and Operation Manual* for your bridge model.
 - ATTO QuickNAV utility (to the computer you are using for setup).
 - c. Go to the ATTO FibreBridge Firmware Download page by clicking **Continue** at the end of the ATTO FibreBridge Description page, and then do the following:
 - Download the bridge firmware file as directed on the download page.

In this step, you are only completing the download portion of the instructions provided in the links. You update the firmware on each bridge later, when instructed to do so in the [Hot-adding the stack of shelves](#) section.

 - Make a copy of the ATTO FibreBridge Firmware Download page and release notes for reference later.
3. Download the latest disk and disk shelf firmware, and make a copy of the installation portion of the instructions for reference later.



In this step, you are only completing the download portion of the instructions provided in the links and making a copy of the installation instructions. You update the firmware on each disk and disk shelf later, when instructed to do so in the [Hot-adding the stack of shelves](#) section.

- a. Download the disk firmware and make a copy of the disk firmware instructions for reference later.

[NetApp Downloads: Disk Drive Firmware](#)

- b. Download the disk shelf firmware and make a copy of the disk shelf firmware instructions for reference later.

[NetApp Downloads: Disk Shelf Firmware](#)

4. Gather the hardware and information needed to use the recommended bridge management interfaces—the ATTO ExpressNAV GUI and ATTO QuickNAV utility:

- a. Acquire a standard Ethernet cable to connect from the bridge Ethernet management 1 port to your network.
- b. Determine a non-default user name and password for accessing the bridges.

It is recommended that you change the default user name and password.

- c. Obtain an IP address, subnet mask, and gateway information for the Ethernet management 1 port on each bridge.
- d. Disable VPN clients on the computer you are using for setup.

Active VPN clients cause the QuickNAV scan for bridges to fail.

5. Acquire four screws for each bridge to flush-mount the bridge “L” brackets securely to the front of the rack.

The openings in the bridge “L” brackets are compliant with rack standard ETA-310-X for 19-inch (482.6 mm) racks.

6. If necessary, update the FC switch zoning to accommodate the new bridges that are being added to the configuration.

If you are using the Reference Configuration Files provided by NetApp, the zones have been created for all ports, so you do not need to make any zoning updates. There must be a storage zone for each switch port that connects to the FC ports of the bridge.

Hot-adding a stack of SAS disk shelves and bridges

You can hot-add a stack of SAS disk shelves and bridges to increase the capacity of the bridges.

The system must meet all of the requirements to hot-add a stack of SAS disk shelves and bridges.

[Preparing to hot-add a stack of SAS disk shelves and bridges](#)

- Hot-adding a stack of SAS disk shelves and bridges is a nondisruptive procedure if all of the interoperability requirements are met.

[NetApp Interoperability Matrix Tool](#)

[Using the Interoperability Matrix Tool to find MetroCluster information](#)

- Multipath HA is the only supported configuration for MetroCluster systems that are using bridges.

Both controller modules must have access through the bridges to the disk shelves in each stack.

- You should hot-add an equal number of disk shelves at each site.
- If you will be using in-band management of the bridge rather than IP management, the steps for configuring the Ethernet port and IP settings can be skipped, as noted in the relevant steps.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.



If you insert a SAS cable into the wrong port, when you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that the cable has been moved to another port.

Steps

1. Properly ground yourself.
2. From the console of either controller module, check whether your system has disk autoassignment enabled:

```
storage disk option show
```

The Auto Assign column indicates whether disk autoassignment is enabled.

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
node_A_1	on	on	on	default
node_A_2	on	on	on	default
2 entries were displayed.				

3. Disable the switch ports for the new stack.
4. If configuring for in-band management, connect a cable from FibreBridge RS-232 serial port to the serial (COM) port on a personal computer.

The serial connection will be used for initial configuration, and then in-band management via ONTAP and the FC ports can be used to monitor and manage the bridge.

5. If configuring for IP management, configure the Ethernet management 1 port for each bridge by following the procedure in section 2.0 of the *ATTO FibreBridge Installation and Operation Manual* for your bridge model.

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

When running QuickNAV to configure an Ethernet management port, only the Ethernet management port that is connected by the Ethernet cable is configured. For example, if you also wanted to configure the Ethernet management 2 port, you would need to connect the Ethernet cable to port 2 and run QuickNAV.

6. Configure the bridge.

If you retrieved the configuration information from the old bridge, use the information to configure the new bridge.

Be sure to make note of the user name and password that you designate.

The *ATTO FibreBridge Installation and Operation Manual* for your bridge model has the most current information on available commands and how to use them.



Do not configure time synchronization on ATTO FibreBridge 7600N or 7500N. The time synchronization for ATTO FibreBridge 7600N or 7500N is set to the cluster time after the bridge is discovered by ONTAP. It is also synchronized periodically once a day. The time zone used is GMT and is not changeable.

- a. If configuring for IP management, configure the IP settings of the bridge.

To set the IP address without the QuickNAV utility, you need to have a serial connection to the FibreBridge.

If using the CLI, you must run the following commands:

```
set ipaddress mp1 ip-address
```

```
set ipsubnetmask mp1 subnet-mask
```

```
set ipgateway mp1 x.x.x.x
```

```
set ipdhcp mp1 disabled
```

```
set ethernetspeed mp1 1000
```

- b. Configure the bridge name.

The bridges should each have a unique name within the MetroCluster configuration.

Example bridge names for one stack group on each site:

- bridge_A_1a
- bridge_A_1b
- bridge_B_1a
- bridge_B_1b

If using the CLI, you must run the following command:

```
set bridgename bridgename
```

- c. If running ONTAP 9.4 or earlier, enable SNMP on the bridge:

```
set SNMP enabled
```

In systems running ONTAP 9.5 or later, in-band management can be used to access the bridge via the FC ports rather than the Ethernet port. Beginning with ONTAP 9.8, only in-band management is supported and SNMP management is deprecated.

7. Configure the bridge FC ports.

- a. Configure the data rate/speed of the bridge FC ports.

The supported FC data rate depends on your model bridge.

- The FibreBridge 7600N bridge supports up to 32, 16, or 8 Gbps.
- The FibreBridge 7500N bridge supports up to 16, 8, or 4 Gbps.



The FCDataRate speed you select is limited to the maximum speed supported by both the bridge and the switch to which the bridge port connects. Cabling distances must not exceed the limitations of the SFPs and other hardware.

If using the CLI, you must run the following command:

```
set FCDataRate port-number port-speed
```

- b. If you are configuring a FibreBridge 7500N bridge, configure the connection mode that the port uses to "ptp".



The FCConnMode setting is not required when configuring a FibreBridge 7600N bridge.

If using the CLI, you must run the following command:

```
set FCConnMode port-number ptp
```

- c. If you are configuring a FibreBridge 7600N or 7500N bridge, you must configure or disable the FC2 port.

- If you are using the second port, you must repeat the previous substeps for the FC2 port.
- If you are not using the second port, then you must disable the port:

```
FCPortDisable port-number
```

- d. If you are configuring a FibreBridge 7600N or 7500N bridge, disable the unused SAS ports:

```
SASPortDisable sas-port
```



SAS ports A through D are enabled by default. You must disable the SAS ports that are not being used. If only SAS port A is used, then SAS ports B, C, and D must be disabled.

8. Secure access to the bridge and save the bridge's configuration.

- a. From the controller prompt check the status of the bridges:

```
storage bridge show
```

The output shows which bridge is not secured.

- b. Check the status of the unsecured bridge's ports:

```
info
```

The output shows the status of Ethernet ports MP1 and MP2.

- c. If Ethernet port MP1 is enabled, run the following command:

```
set EthernetPort mp1 disabled
```



If Ethernet port MP2 is also enabled, repeat the previous substep for port MP2.

d. Save the bridge's configuration.

You must run the following commands:

```
SaveConfiguration
```

```
FirmwareRestart
```

You are prompted to restart the bridge.

9. Update the FibreBridge firmware on each bridge.

If the new bridge is the same type as the partner bridge upgrade to the same firmware as the partner bridge. If the new bridge is a different type to the partner bridge, upgrade to the latest firmware supported by the bridge and version of ONTAP. See the section "Updating firmware on a FibreBridge bridge" in *MetroCluster Maintenance*.

10. Cable the disk shelves to the bridges:

a. Daisy-chain the disk shelves in each stack.

The *Installation Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

b. For each stack of disk shelves, cable IOM A of the first shelf to SAS port A on FibreBridge A, and then cable IOM B of the last shelf to SAS port A on FibreBridge B.

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

Each bridge has one path to its stack of disk shelves; bridge A connects to the A-side of the stack through the first shelf, and bridge B connects to the B-side of the stack through the last shelf.



The bridge SAS port B is disabled.

11. Verify that each bridge can detect all of the disk drives and disk shelves to which the bridge is connected.

If you are using the...	Then...
-------------------------	---------

ATTO ExpressNAV GUI	<p>a. In a supported web browser, enter the IP address of a bridge in the browser box.</p> <p>You are brought to the ATTO FibreBridge home page, which has a link.</p> <p>b. Click the link, and then enter your user name and the password that you designated when you configured the bridge.</p> <p>The ATTO FibreBridge status page appears with a menu to the left.</p> <p>c. Click Advanced in the menu.</p> <p>d. View the connected devices:</p> <pre>sastargets</pre> <p>e. Click Submit.</p>
Serial port connection	<p>View the connected devices:</p> <pre>sastargets</pre>

The output shows the devices (disks and disk shelves) to which the bridge is connected. The output lines are sequentially numbered so that you can quickly count the devices.



If the text response truncated appears at the beginning of the output, you can use Telnet to connect to the bridge, and then view all of the output by using the `sastargets` command.

The following output shows that 10 disks are connected:

Tgt	VendorID	ProductID	Type	SerialNumber
0	NETAPP	X410_S15K6288A15	DISK	3QP1CLE300009940UHJV
1	NETAPP	X410_S15K6288A15	DISK	3QP1ELF600009940V1BV
2	NETAPP	X410_S15K6288A15	DISK	3QP1G3EW00009940U2M0
3	NETAPP	X410_S15K6288A15	DISK	3QP1EWMP00009940U1X5
4	NETAPP	X410_S15K6288A15	DISK	3QP1FZLE00009940G8YU
5	NETAPP	X410_S15K6288A15	DISK	3QP1FZLF00009940TZKZ
6	NETAPP	X410_S15K6288A15	DISK	3QP1CEB400009939MGXL
7	NETAPP	X410_S15K6288A15	DISK	3QP1G7A900009939FNNT
8	NETAPP	X410_S15K6288A15	DISK	3QP1FY0T00009940G8PA
9	NETAPP	X410_S15K6288A15	DISK	3QP1FXW600009940VERQ

12. Verify that the command output shows that the bridge is connected to all of the appropriate disks and disk shelves in the stack.

If the output is...	Then...
Correct	Repeat Step 11 for each remaining bridge.

Not correct	<p>a. Check for loose SAS cables or correct the SAS cabling by repeating Step 10.</p> <p>b. Repeat Step 11.</p>
-------------	---

13. If you are configuring a fabric-attached MetroCluster configuration, cable each bridge to the local FC switches, using the cabling shown in the table for your configuration, switch model, and FC-to-SAS bridge model:



Brocade and Cisco switches use different port numbering, as shown in the following tables.

- On Brocade switches, the first port is numbered “0”.
- On Cisco switches, the first port is numbered “1”.

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)												
DR GROUP 1												
			Brocade 6505		Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G620, Brocade G620-1, Brocade G630, Brocade G630-1		Brocade G720	
Component	Port	Port	Switc h 1	Switc h 2	Switc h 1	Switc h 2	Switc h 1	Switc h 2	Switc h 1	Switc h 2	Switc h 1	Switc h 2
Stack 1	bridge _x_1a	FC1	8		8		8		8		10	
		FC2	-	8	-	8	-	8	-	8	-	10
	bridge _x_1B	FC1	9	-	9	-	9	-	9	-	11	-
		FC2	-	9	-	9	-	9	-	9	-	11
Stack 2	bridge _x_2a	FC1	10	-	10	-	10	-	10	-	14	-
		FC2	-	10	-	10	-	10	-	10	-	14
	bridge _x_2B	FC1	11	-	11	-	11	-	11	-	17	-
		FC2	-	11	-	11	-	11	-	11	-	17

Stack 3	bridge_x_3a	FC1	12	-	12	-	12	-	12	-	18	-
		FC2	-	12	-	12	-	12	-	12	-	18
	bridge_x_3B	FC1	13	-	13	-	13	-	13	-	19	-
		FC2	-	13	-	13	-	13	-	13	-	19
Stack y	bridge_x_ya	FC1	14	-	14	-	14	-	14	-	20	-
		FC2	-	14	-	14	-	14	-	14	-	20
	bridge_x_yb	FC1	15	-	15	-	15	-	15	-	21	-
		FC2	-	15	-	15	-	15	-	15	-	21
 Additional bridges can be cabled to ports 16, 17, 20 and 21 in G620, G630, G620-1, and G630-1 switches.												

Configurations using FibreBridge 7500N or 7600N using both FC ports (FC1 and FC2)										
DR GROUP 2										
			Brocade G620, Brocade G620-1, Brocade G630, Brocade G630-1		Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G720	
Component		Port	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	switch 2
Stack 1	bridge_x_51a	FC1	26	-	32	-	56	-	32	-
		FC2	-	26	-	32	-	56	-	32
	bridge_x_51b	FC1	27	-	33	-	57	-	33	-
		FC2	-	27	-	33	-	57	-	33
Stack 2	bridge_x_52a	FC1	30	-	34	-	58	-	34	-
		FC2	-	30	-	34	-	58	-	34
	bridge_x_52b	FC1	31	-	35	-	59	-	35	-
		FC2	-	31	-	35	-	59	-	35

Stack 3	bridge_x_53a	FC1	32	-	36	-	60	-	36	-	
		FC2	-	32	-	36	-	60	-	36	
	bridge_x_53b	FC1	33	-	37	-	61	-	37	-	
		FC2	-	33	-	37	-	61	-	37	
Stack y	bridge_x_5ya	FC1	34	-	38	-	62	-	38	-	
		FC2	-	34	-	38	-	62	-	38	
	bridge_x_5yb	FC1	35	-	39	-	63	-	39	-	
		FC2	-	35	-	39	-	63	-	39	
		Additional bridges can be cabled to ports 36 - 39 in G620, G630, G620-1, and G-630-1 switches.									

Configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only

DR GROUP 1

		Brocade 6505		Brocade 6510, Brocade DCX 8510-8		Brocade 6520		Brocade G620, brocade G620-1, Brocade G630, Brocade G630-1		Brocade G720	
Comp onent	Port	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2	Switch 1	Switch 2
Stack 1	bridge_x_1a	8		8		8		8		10	
	bridge_x_1b	-	8	-	8	-	8	-	8	-	10
Stack 2	bridge_x_2a	9	-	9	-	9	-	9	-	11	-
	bridge_x_2b	-	9	-	9	-	9	-	9	-	11

Stack 3	bridge_x_3a	10	-	10	-	10	-	10	-	14	-
	bridge_x_4b	-	10	-	10	-	10	-	10	-	14
Stack y	bridge_x_ya	11	-	11	-	11	-	11	-	15	-
	bridge_x_yb	-	11	-	11	-	11	-	11	-	15



Additional bridges can be cabled to ports 12 - 17, 20 and 21 in G620, G630, G620-1, and G630-1 switches. Additional bridges can be cabled to ports 16 - 17, 20 and 21 G720 switches.

Configurations using FibreBridge 7500N or 7600N using one FC port (FC1 or FC2) only

DR GROUP 2

		Brocade G720		Brocade G620, Brocade G620-1, Brocade G630, Brocade G630-1		Brocade 6510, Brocade DCX 8510-8		Brocade 6520	
Stack 1	bridge_x_51a	32	-	26	-	32	-	56	-
	bridge_x_51b	-	32	-	26	-	32	-	56
Stack 2	bridge_x_52a	33	-	27	-	33	-	57	-
	bridge_x_52b	-	33	-	27	-	33	-	57
Stack 3	bridge_x_53a	34	-	30	-	34	-	58	-
	bridge_x_54b	-	34	-	30	-	34	-	58
Stack y	bridge_x_ya	35	-	31	-	35	-	59	-
	bridge_x_yb	-	35	-	31	-	35	-	59



Additional bridges can be cabled to ports 32 through 39 in G620, G630, G620-1, and G630-1 switches. Additional bridges can be cabled to ports 36 through 39 in G720 switches.

14. If you are configuring a bridge-attached MetroCluster system, cable each bridge to the controller modules:
 - a. Cable FC port 1 of the bridge to a 16 Gb or 8 Gb FC port on the controller module in cluster_A.
 - b. Cable FC port 2 of the bridge to the same speed FC port of the controller module in cluster_A.
 - c. Repeat these substeps on other subsequent bridges until all of the bridges have been cabled.
15. Update the disk drive firmware to the most current version from the system console:

```
disk_fw_update
```

You must run this command on both controller modules.

[NetApp Downloads: Disk Drive Firmware](#)

16. Update the disk shelf firmware to the most current version by using the instructions for the downloaded firmware.

You can run the commands in the procedure from the system console of either controller module.

[NetApp Downloads: Disk Shelf Firmware](#)

17. If your system does not have disk autoassignment enabled, assign disk drive ownership.

[Disk and aggregate management](#)



If you are splitting the ownership of a single stack of disk shelves among multiple controller modules, you must disable disk autoassignment on both nodes in the cluster (`storage disk option modify -autoassign off *`) before assigning disk ownership; otherwise, when you assign any single disk drive, the remaining disk drives might be automatically assigned to the same controller module and pool.



You must not add disk drives to aggregates or volumes until after the disk drive firmware and disk shelf firmware have been updated and the verification steps in this task have been completed.

18. Enable the switch ports for the new stack.
19. Verify the operation of the MetroCluster configuration in ONTAP:
 - a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

20. If applicable, repeat this procedure for the partner site.

Related information

[In-band management of the FC-to-SAS bridges](#)

Hot add a SAS disk shelf to a stack of SAS disk shelves

You can hot-add a disk shelf when you want to increase storage without any reduction in performance.

Step 1: Prepare to hot-add a SAS disk shelf

To prepare for hot-adding a SAS disk shelf, you need to download documents along with the disk drive and disk shelf firmware.

Before you begin

- Verify that your system is a supported configuration and is running a supported version of ONTAP.
- Verify that all disk drives and disk shelves in the system are running the latest firmware version.

You might want to update the disk and shelf firmware throughout the MetroCluster configuration before you add shelves.

[Upgrade, revert, or downgrade](#)

Steps

1. Download or view the following documents from the NetApp Support Site:
 - [Interoperability Matrix Tool](#)
 - The *Installation Guide* for your disk shelf model.
2. Verify that the disk shelf you are hot-adding is supported.

[Interoperability Matrix Tool](#)

3. Download the latest disk and disk shelf firmware:



In this step, you only complete the download portion of the instructions. You need to follow the steps in [hot-add a disk shelf](#) to install the disk shelf.

- a. Download the disk firmware and make a copy of the disk firmware instructions for reference later.

[NetApp Downloads: Disk Drive Firmware](#)

- b. Download the disk shelf firmware and make a copy of the disk shelf firmware instructions for reference later.

[NetApp Downloads: Disk Shelf Firmware](#)

Step 2: Hot-add a disk shelf

Use the following procedure to hot-add a disk shelf to a stack.

Before you begin

- Verify that the system meets all of the requirements in [Prepare to hot-add SAS disk shelves](#).
- Verify that your environment meets one of the following scenarios before you hot-add a shelf:
 - You have two FibreBridge 7500N bridges connected to a stack of SAS disk shelves.
 - You have two FibreBridge 7600N bridges connected to a stack of SAS disk shelves.
 - You have one FibreBridge 7500N bridge and one FibreBridge 7600N bridge connected to a stack of SAS disk shelves.

About this task

- This procedure is for hot-adding a disk shelf to the last disk shelf in a stack.

This procedure is written with the assumption that the last disk shelf in a stack is connected from IOM A to bridge A and from IOM B to bridge B.

- This is a nondisruptive procedure.
- You should hot-add an equal number of disk shelves at each site.
- If you are hot-adding more than one disk shelf, you must hot-add one disk shelf at a time.

Each pair of FibreBridge 7500N or 7600N bridges can support up to four stacks.



Hot-adding a disk shelf requires you to update the disk drive firmware on the hot-added disk shelf by running the `storage disk firmware update` command in advanced mode. Running this command can be disruptive if the firmware on existing disk drives in your system is an older version.

If you insert a SAS cable into the wrong port, after you remove the cable from a SAS port, you must wait at least 120 seconds before plugging the cable into a different SAS port. If you fail to do so, the system will not recognize that you have moved the cable to a different port.

Steps

1. Properly ground yourself.

2. Verify disk shelf connectivity from the system console of either controller:

```
sysconfig -v
```

The output is similar to the following:

- Each bridge on a separate line and under each FC port to which it is visible; for example, hot-adding a disk shelf to a set of FibreBridge 7500N bridges results in the following output:

```
FC-to-SAS Bridge:
cisco_A_1-1:9.126L0: ATTO  FibreBridge7500N 2.10  FB7500N100189
cisco_A_1-2:1.126L0: ATTO  FibreBridge7500N 2.10  FB7500N100162
```

- Each disk shelf on a separate line under each FC port to which it is visible:

```
Shelf    0: IOM6  Firmware rev. IOM6 A: 0173 IOM6 B: 0173
Shelf    1: IOM6  Firmware rev. IOM6 A: 0173 IOM6 B: 0173
```

- Each disk drive on a separate line under each FC port to which it is visible:

```
cisco_A_1-1:9.126L1    : NETAPP    X421_HCOBD450A10 NA01 418.0GB
(879097968 520B/sect)
cisco_A_1-1:9.126L2    : NETAPP    X421_HCOBD450A10 NA01 418.0GB
(879097968 520B/sect)
```

3. Check whether your system has disk auto-assignment enabled from the console of either controller:

```
storage disk option show
```

The auto-assignment policy is shown in the Auto Assign column.

Node	BKg. FW. Upd.	Auto Copy	Auto Assign	Auto Assign Policy
node_A_1	on	on	on	default
node_A_2	on	on	on	default

2 entries were displayed.

4. If your system does not have disk auto-assignment enabled, or if disk drives in the same stack are owned by both controllers, assign disk drives to the appropriate pools.

[Disk and aggregate management](#)



- If you are splitting a single stack of disk shelves between two controllers, disk auto-assignment must be disabled before you assign disk ownership; otherwise, when you assign any single disk drive, the remaining disk drives might be automatically assigned to the same controller and pool.

The `storage disk option modify -node <node-name> -autoassign off` command disables disk autoassignment.

- You cannot add drives to aggregates or volumes until after you have updated the disk drive and disk shelf firmware.

5. Update the disk shelf firmware to the most current version by using the instructions for the downloaded firmware.

You can run the commands in the procedure from the system console of either controller.

[NetApp Downloads: Disk Shelf Firmware](#)

6. Install and cable the disk shelf:



Do not force a connector into a port. The mini-SAS cables are keyed; when oriented correctly into a SAS port, the SAS cable clicks into place and the disk shelf SAS port LNK LED illuminates green. For disk shelves, you insert a SAS cable connector with the pull tab oriented up (on the topside of the connector).

- a. Install the disk shelf, power it on, and set the shelf ID.

The *Installation Guide* for your disk shelf model provides detailed information about installing disk shelves.



You must power-cycle the disk shelf and keep the shelf IDs unique for each SAS disk shelf within the entire storage system.

- b. Disconnect the SAS cable from the IOM B port of the last shelf in the stack, and then reconnect it to the same port in the new shelf.

The other end of this cable remains connected to bridge B.

- c. Daisy-chain the new disk shelf by cabling the new shelf IOM ports (of IOM A and IOM B) to the last shelf IOM ports (of IOM A and IOM B).

The *Installation Guide* for your disk shelf model provides detailed information about daisy-chaining disk shelves.

7. Update the disk drive firmware to the most current version from the system console.

[NetApp Downloads: Disk Drive Firmware](#)

- a. Change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with **y** when prompted to continue into advanced mode and see the advanced mode prompt (***>**).

- b. Update the disk drive firmware to the most current version from the system console:

```
storage disk firmware update
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

- d. Repeat the previous substeps on the other controller.

8. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node <node-name> sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Active IQ Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

9. If you are hot-adding more than one disk shelf, repeat the previous steps for each disk shelf that you are hot-adding.

Hot-removing storage from a MetroCluster FC configuration

You can hot-remove drive shelves—physically remove shelves that have had the aggregates removed from the drives—from a MetroCluster FC configuration that is up and serving data. You can hot-remove one or more shelves from anywhere within a stack of shelves or remove a stack of shelves.

- Your system must be a multipath HA, multipath, quad-path HA, or quad-path configuration.
- In a four-node MetroCluster FC configuration, the local HA pair cannot be in a takeover state.

- You must have already removed all aggregates from the drives in the shelves that you are removing.



If you attempt this procedure on non-MetroCluster FC configurations with aggregates on the shelf you are removing, you could cause the system to fail with a multidrive panic.

Removing aggregates involves splitting the mirrored aggregates on the shelves you are removing, and then re-creating the mirrored aggregates with another set of drives.

Disk and aggregate management

- You must have removed drive ownership after removing the aggregates from the drives in the shelves that you are removing.

Disk and aggregate management

- If you are removing one or more shelves from within a stack, you must have factored the distance to bypass the shelves that you are removing.

If the current cables are not long enough, you need to have longer cables available.

This task applies to the following MetroCluster FC configurations:

- Direct-attached MetroCluster FC configurations, in which the storage shelves are directly connected to the storage controllers with SAS cables
- Fabric-attached or bridge-attached MetroCluster FC configurations, in which the storage shelves are connected using FC-to-SAS bridges

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

2. Set the privilege level to advanced:

```
set -privilege advanced
```

3. Verify that no mailbox drive is on the shelves:

```
storage failover mailbox-disk show
```

4. Remove the shelf according to the steps for the relevant scenario.

Scenario	Steps
To remove an aggregate when the shelf contains either unmirrored, mirrored, or both types of aggregate...	<p>a. Use the <code>storage aggregate delete -aggregate <i>aggregate name</i></code> command to remove the aggregate.</p> <p>b. Use the standard procedure to remove ownership of all drives in that shelf, and then physically remove the shelf.</p> <p>Follow the instructions in the <i>SAS Disk Shelves Service Guide</i> for your shelf model to hot-remove shelves.</p>

To remove a plex from a mirrored aggregate, you need to unmirror the aggregate.

- a. Identify the plex that you want to remove by using the `run -node local sysconfig -r` command.

In the following example, you can identify the plex from the line `Plex`

`/dpg_mcc_8020_13_a1_aggr1/plex0`. In this case, the plex to specify is "plex0".

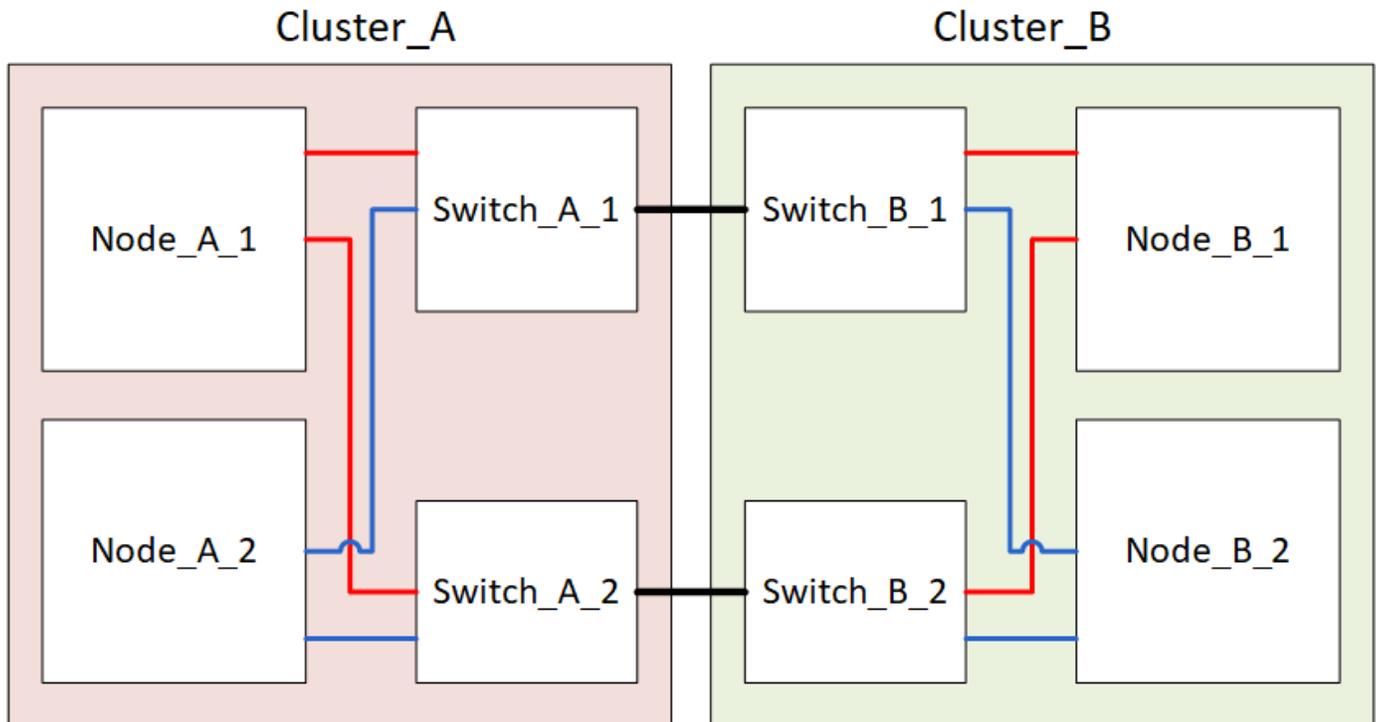
```
dpgmcc_8020_13_a1a2::storage
aggregate> run -node local
sysconfig -r
*** This system has taken over
dpg-mcc-8020-13-a1
Aggregate
dpg_mcc_8020_13_a1_aggr1
(online, raid_dp, mirrored)
(block checksums)
Plex
/dpg_mcc_8020_13_a1_aggr1/plex
0 (online, normal, active,
pool0)
RAID group
/dpg_mcc_8020_13_a1_aggr1/plex
0/rg0 (normal, block
checksums)
RAID Disk Device
HA  SHELF BAY CHAN Pool Type
RPM  Used (MB/blks)  Phys
(MB/blks)
-----
-----
-----
-----
dparity  mcc-cisco-8Gb-
fab-2:1-1.126L16 0c 32 15
FC:B 0 SAS 15000
272000/557056000
274845/562884296
parity  mcc-cisco-8Gb-
fab-2:1-1.126L18 0c 32 17
FC:B 0 SAS 15000
272000/557056000
274845/562884296
data  mcc-cisco-8Gb-
fab-2:1-1.126L19 0c 32 18
FC:B 0 SAS 15000
272000/557056000
274845/562884296
data  mcc-cisco-8Gb-
```

Power off and power on a single site in a MetroCluster FC configuration

If you need to perform site maintenance or relocate a single site in a MetroCluster FC configuration, you must know how to power off and power on the site.

If you need to relocate and reconfigure a site (for example, if you need to expand from a four-node to an eight-node cluster), you cannot complete these tasks at the same time. This procedure only covers the steps that are required to perform site maintenance or to relocate a site without changing its configuration.

The following diagram shows a MetroCluster configuration. Cluster_B is powered off for maintenance.



Power off a MetroCluster site

You must power off a site and all of the equipment before site maintenance or relocation can begin.

About this task

All the commands in the following steps are issued from the site that remains powered on.

Steps

- Before you begin, check that any non-mirrored aggregates at the site are offline.

- Verify the operation of the MetroCluster configuration in ONTAP:
 - Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- Check for any health alerts on both clusters:

```
system health alert show
```

- Confirm the MetroCluster configuration and that the operational mode is normal:

```

FC:B 0 SAS 15000
272000/557056000
274845/562884296
data mcc-cisco-8Gb-
fab-3:1-1.126L22 0d 32 21
272000/557056000
274845/562884296

```

```

fab-3:1-1.126L37 0d 34 10
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

parity mcc-cisco-8Gb-
fab-3:1-1.126L14 0d 33 13
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

data mcc-cisco-8Gb-
fab-3:1-1.126L41 0d 34 14
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

data mcc-cisco-8Gb-
fab-3:1-1.126L15 0d 33 14
FC:A 1 SAS 15000
272000/557056000
280104/573653840

```

```

data mcc-cisco-8Gb-
fab-3:1-1.126L45 0d 34 18

```

```
metrocluster show
```

```
FC:A 1 SAS 15000  
272000/557056000  
280104/573653840
```

d. Perform a MetroCluster check:

```
metrocluster check run
```

e. Display the results of the MetroCluster check:

```
metrocluster check show
```

f. Check for any health alerts on the switches (if present):

```
storage switch show
```

g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

h. After running Config Advisor, review the tool's output and address any issues discovered.

b. Use the storage aggregate plex delete -aggregate *aggr_name* -plex *plex_name* command to remove the plex.

plex defines the plex name, such as "plex3" or "plex6".

c. Use the standard procedure to remove ownership of all drives in that shelf, and then physically remove the shelf.

Follow the instructions in the *SAS Disk Shelves Service Guide* for your shelf model to hot remove shelves.

3. From the site you want to remain up, implement the switchover:

```
metrocluster switchover
```

```
cluster_A::*> metrocluster switchover
```

The operation can take several minutes to complete.

The unmirrored aggregates will only be online after a switchover if the remote disks in the aggregate are accessible. If the ISLs fail, the local node might be unable to access the data in the unmirrored remote disks. The failure of an aggregate can lead to a reboot of the local node.

4. Monitor and verify the completion of the switchover:

```
metrocluster operation show
```

```
cluster_A::*> metrocluster operation show
  Operation: Switchover
  Start time: 10/4/2012 19:04:13
  State: in-progress
  End time: -
  Errors:
```

```
cluster_A::*> metrocluster operation show
  Operation: Switchover
  Start time: 10/4/2012 19:04:13
  State: successful
  End time: 10/4/2012 19:04:22
  Errors: -
```

5. Move any volumes and LUNs that belong to unmirrored aggregates offline.

a. Move the volumes offline.

```
cluster_A::* volume offline <volume name>
```

b. Move the LUNs offline.

```
cluster_A::* lun offline lun_path <lun_path>
```

6. Move unmirrored aggregates offline: storage aggregate offline

```
cluster_A*::> storage aggregate offline -aggregate <aggregate-name>
```

7. Depending on your configuration and ONTAP version, identify and move offline affected plexes that are located at the disaster site (Cluster_B).

You should move the following plexes offline:

- Non-mirrored plexes residing on disks located at the disaster site.

If you do not move the non-mirrored plexes at the disaster site offline, an outage might occur when the disaster site is later powered off.

- Mirrored plexes residing on disks located at the disaster site for aggregate mirroring. After they are moved offline, the plexes are inaccessible.

a. Identify the affected plexes.

Plexes that are owned by nodes at the surviving site consist of Pool1 disks. Plexes that are owned by nodes at the disaster site consist of Pool0 disks.

```

Cluster_A::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate      plex  status          is-online pool
-----
Node_B_1_aggr0 plex0 normal,active true      0
Node_B_1_aggr0 plex1 normal,active true      1

Node_B_2_aggr0 plex0 normal,active true      0
Node_B_2_aggr0 plex5 normal,active true      1

Node_B_1_aggr1 plex0 normal,active true      0
Node_B_1_aggr1 plex3 normal,active true      1

Node_B_2_aggr1 plex0 normal,active true      0
Node_B_2_aggr1 plex1 normal,active true      1

Node_A_1_aggr0 plex0 normal,active true      0
Node_A_1_aggr0 plex4 normal,active true      1

Node_A_1_aggr1 plex0 normal,active true      0
Node_A_1_aggr1 plex1 normal,active true      1

Node_A_2_aggr0 plex0 normal,active true      0
Node_A_2_aggr0 plex4 normal,active true      1

Node_A_2_aggr1 plex0 normal,active true      0
Node_A_2_aggr1 plex1 normal,active true      1
14 entries were displayed.

Cluster_A::>

```

The affected plexes are those that are remote to cluster A. The following table shows whether the disks are local or remote relative to cluster A:

Node	Disks in pool	Should the disks be set offline?	Example of plexes to be moved offline
Node_A_1 and Node_A_2	Disks in pool 0	No. Disks are local to cluster A.	-
	Disks in pool 1	Yes. Disks are remote to cluster A.	Node_A_1_aggr0/plex4 Node_A_1_aggr1/plex1 Node_A_2_aggr0/plex4 Node_A_2_aggr1/plex1

Node_B_1 and Node_B_2	Disks in pool 0	Yes. Disks are remote to cluster A.	Node_B_1_aggr1/plex0 Node_B_1_aggr0/plex0 Node_B_2_aggr0/plex0 Node_B_2_aggr1/plex0
	Disks in pool 1	No. Disks are local to cluster A.	-

b. Move the affected plexes offline:

```
storage aggregate plex offline
```

```
storage aggregate plex offline -aggregate Node_B_1_aggr0 -plex plex0
```



Perform this step for all plexes that have disks that are remote to Cluster_A.

8. Persistently offline the ISL switch ports according to the switch type.

Switch type	Action
-------------	--------

For Brocade FC switches...

- a. Use the `portcfgpersistentdisable <port>` command to persistently disable the ports as shown in the following example. This must be done on both switches at the surviving site.

```
Switch_A_1:admin> portcfgpersistentdisable 14
Switch_A_1:admin> portcfgpersistentdisable 15
Switch_A_1:admin>
```

- b. Verify that the ports are disabled using the `switchshow` command shown in the following example:

```
Switch_A_1:admin> switchshow
switchName:      Switch_A_1
switchType:      109.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:    2
switchId:        fffc02
switchWwn:       10:00:00:05:33:88:9c:68
zoning:          ON (T5_T6)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
Address Mode:    0

  Index Port Address Media Speed State      Proto
  =====
  ...
  14  14  020e00  id   16G  No_Light  FC
Disabled (Persistent)
  15  15  020f00  id   16G  No_Light  FC
Disabled (Persistent)
  ...
Switch_A_1:admin>
```

For Cisco FC switches...

- a. Use the `interface` command to persistently disable the ports. The following example shows ports 14 and 15 being disabled:

```
Switch_A_1# conf t
Switch_A_1(config)# interface fc1/14-15
Switch_A_1(config)# shut

Switch_A_1(config-if)# end
Switch_A_1# copy running-config startup-config
```

- b. Verify that the switch port is disabled using the `show interface brief` command as shown in the following example:

```
Switch_A_1# show interface brief
Switch_A_1
```

9. Power off the equipment at the disaster site.

The following equipment must be turned off in the order shown:

- Storage controllers - the storage controllers should currently be at the `LOADER` prompt, you must power them off completely.
- MetroCluster FC switches
- ATTO FibreBridges (if present)
- Storage shelves

Relocating the powered-off site of the MetroCluster

After the site is powered off, you can begin maintenance work. The procedure is the same whether the MetroCluster components are relocated within the same data center or relocated to a different data center.

- The hardware should be cabled in the same way as the previous site.
- If the Inter-Switch Link (ISL) speed, length, or number has changed, they all need to be reconfigured.

Steps

1. Verify that the cabling for all components is carefully recorded so that it can be correctly reconnected at the new location.
2. Physically relocate all the hardware, storage controllers, FC switches, FibreBridges, and storage shelves.
3. Configure the ISL ports and verify the intersite connectivity.
 - a. Power on the FC switches.



Do **not** power up any other equipment.

- b. Enable the ports.

Enable the ports according to the correct switch types in the following table:

Switch type	Command
-------------	---------

For Brocade FC switches...

- a. Use the `portcfgpersistentenable <port number>` command to persistently enable the port. This must be done on both switches at the surviving site.

The following example shows ports 14 and 15 being enabled on Switch_A_1.

```
switch_A_1:admin>
portcfgpersistentenable 14
switch_A_1:admin>
portcfgpersistentenable 15
switch_A_1:admin>
```

- b. Verify that the switch port is enabled: `switchshow`

The following example shows that ports 14 and 15 are enabled:

```
switch_A_1:admin> switchshow
switchName: Switch_A_1
switchType: 109.1

switchState:    Online
switchMode: Native
switchRole: Principal
switchDomain:    2
switchId:    fffc02
switchWwn:    10:00:00:05:33:88:9c:68
zoning:    ON (T5_T6)
switchBeacon:    OFF
FC Router:    OFF
FC Router BB Fabric ID: 128
Address Mode:    0

Index Port Address Media Speed State
Proto
=====
=====
...
 14 14 020e00 id 16G Online
FC E-Port 10:00:00:05:33:86:89:cb
"Switch_A_1"
 15 15 020f00 id 16G Online
FC E-Port 10:00:00:05:33:86:89:cb
"Switch_A_1" (downstream)
...
switch_A_1:admin>
```

For Cisco FC switches...

a. Enter the `interface` command to enable the port.

The following example shows ports 14 and 15 being enabled on Switch_A_1.

```
switch_A_1# conf t
switch_A_1(config)# interface fc1/14-15
switch_A_1(config)# no shut
switch_A_1(config-if)# end
switch_A_1# copy running-config
startup-config
```

b. Verify that the switch port is enabled: `show interface brief`

```
switch_A_1# show interface brief
switch_A_1#
```

4. Use tools on the switches (as they are available) to verify the intersite connectivity.



You should only proceed if the links are correctly configured and stable.

5. Disable the links again if they are found to be stable.

Disable the ports based on whether you are using Brocade or Cisco switches as shown in the following table:

Switch type	Command
-------------	---------

For Brocade FC switches...

- a. Enter the `portcfgpersistentdisable <port_number>` command to persistently disable the port.

This must be done on both switches at the surviving site. The following example shows ports 14 and 15 being disabled on Switch_A_1:

```
switch_A_1:admin> portpersistentdisable
14
switch_A_1:admin> portpersistentdisable
15
switch_A_1:admin>
```

- b. Verify that the switch port is disabled: `switchshow`

The following example shows that ports 14 and 15 are disabled:

```
switch_A_1:admin> switchshow
switchName: Switch_A_1
switchType: 109.1
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 2
switchId: fffc02
switchWwn: 10:00:00:05:33:88:9c:68
zoning: ON (T5_T6)
switchBeacon: OFF
FC Router: OFF
FC Router BB Fabric ID: 128
Address Mode: 0

Index Port Address Media Speed State
Proto
=====
=====
...
14 14 020e00 id 16G No_Light
FC Disabled (Persistent)
15 15 020f00 id 16G No_Light
FC Disabled (Persistent)
...
switch_A_1:admin>
```

For Cisco FC switches...

a. Disable the port using the `interface` command.

The following example shows ports fc1/14 and fc1/15 being disabled on Switch A_1:

```
switch_A_1# conf t

switch_A_1(config)# interface fc1/14-15
switch_A_1(config)# shut
switch_A_1(config-if)# end
switch_A_1# copy running-config startup-
config
```

b. Verify that the switch port is disabled using the `show interface brief` command.

```
switch_A_1# show interface brief
switch_A_1#
```

Powering on the MetroCluster configuration and returning to normal operation

After maintenance has been completed or the site has been moved, you must power on the site and reestablish the MetroCluster configuration.

About this task

All the commands in the following steps are issued from the site that you power on.

Steps

1. Power on the switches.

You should power on the switches first. They might have been powered on during the previous step if the site was relocated.

- a. Reconfigure the Inter-Switch Link (ISL) if required or if this was not completed as part of the relocation.
 - b. Enable the ISL if fencing was completed.
 - c. Verify the ISL.
2. Disable the ISLs on the FC switches.
 3. Power on the shelves and allow enough time for them to power on completely.
 4. Power on the FibreBridge bridges.
 - a. On the FC switches, verify that the ports connecting the bridges are coming online.

You can use a command such as `switchshow` for Brocade switches, and `show interface brief` for Cisco switches.

b. Verify that the shelves and disks on the bridges are clearly visible.

You can use a command such as `sastargets` on the ATTO CLI.

5. Enable the ISLs on the FC switches.

Enable the ports based on whether you are using Brocade or Cisco switches as shown in the following table:

Switch type	Command
-------------	---------

For Brocade FC switches...

- a. Enter the `portcfgpersistentenable <port>` command to persistently enable the ports. This must be done on both switches at the surviving site.

The following example shows ports 14 and 15 being enabled on Switch_A_1:

```
Switch_A_1:admin> portcfgpersistentenable 14
Switch_A_1:admin> portcfgpersistentenable 15
Switch_A_1:admin>
```

- b. Verify that the switch port is enabled using the

`switchshow` command:

```
switch_A_1:admin> switchshow
switchName:      Switch_A_1
switchType:      109.1
switchState:     Online
switchMode:      Native
switchRole:      Principal
switchDomain:     2
switchId:        fffc02
switchWwn:       10:00:00:05:33:88:9c:68
zoning:          ON (T5_T6)
switchBeacon:    OFF
FC Router:       OFF
FC Router BB Fabric ID: 128
Address Mode:    0

  Index Port Address Media Speed State      Proto
  =====
  ...
  14  14   020e00   id   16G   Online   FC
E-Port 10:00:00:05:33:86:89:cb "Switch_A_1"
  15  15   020f00   id   16G   Online   FC
E-Port 10:00:00:05:33:86:89:cb "Switch_A_1"
(downstream)
  ...
switch_A_1:admin>
```

For Cisco FC switches...

- a. Use the `interface` command to enable the ports.

The following example shows port fc1/14 and fc1/15 being enabled on Switch A_1:

```
switch_A_1# conf t
switch_A_1(config)# interface fc1/14-15
switch_A_1(config)# no shut
switch_A_1(config-if)# end
switch_A_1# copy running-config startup-config
```

- b. Verify that the switch port is disabled:

```
switch_A_1# show interface brief
switch_A_1#
```

6. Verify that the storage is visible.

- a. Verify that the storage is visible from the surviving site. Bring the offline plexes back online to restart the resync operation and reestablish the SyncMirror.
- b. Verify that the local storage is visible from the node in Maintenance mode:

```
disk show -v
```

7. Reestablish the MetroCluster configuration.

Follow the instructions in [Verifying that your system is ready for a switchback](#) to perform healing and switchback operations according to your MetroCluster configuration.

Powering off an entire MetroCluster FC configuration

You must power off the entire MetroCluster FC configuration and all of the equipment before site maintenance or relocation can begin.

About this task

You must perform the steps in this procedure from both sites, at the same time.



Beginning with ONTAP 9.8, the **storage switch** command is replaced with **system switch**. The following steps show the **storage switch** command, but if you are running ONTAP 9.8 or later, the **system switch** command is preferred.

Steps

1. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.
 - a. Confirm the MetroCluster configuration and that the operational mode is normal.

```
metrocluster show
```

- b. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

- c. Run the following command:

```
storage bridge show
```

- d. Run the following command:

```
storage port show
```

- e. Run the following command:

```
storage switch show
```

- f. Run the following command:

```
network port show
```

- g. Perform a MetroCluster check:

```
metrocluster check run
```

- h. Display the results of the MetroCluster check:

```
metrocluster check show
```

2. Disable AUSO by modifying the AUSO Failure Domain to

```
auso-disabled
```

```
cluster_A_site_A::*>metrocluster modify -auto-switchover-failure-domain  
auso-disabled
```

3. Verify the change using the command

```
metrocluster operation show
```

```
cluster_A_site_A::*> metrocluster operation show  
Operation: modify  
State: successful  
Start Time: 4/25/2020 20:20:36  
End Time: 4/25/2020 20:20:36  
Errors: -
```

4. Halt the nodes by using the following command:

```
halt
```

- For a four-node or eight-node MetroCluster configuration, use the **inhibit-takeover** and **skip-**

lif-migration-before-shutdown parameters:

```
system node halt -node node1_SiteA -inhibit-takeover true -ignore
-quorum-warnings true -skip-lif-migration-before-shutdown true
```

- For a two-node MetroCluster configuration, use the command:

```
system node halt -node node1_SiteA -ignore-quorum-warnings true
```

5. Power off the following equipment at the site:

- Storage controllers
- MetroCluster FC switches (if in use and the configuration is not a a two-node stretch configuration)
- ATTO FibreBridges
- Storage shelves

6. Wait for thirty minutes and then power on the following equipment at the site:

- Storage shelves
- ATTO FibreBridges
- MetroCluster FC switches
- Storage controllers

7. After the controllers are powered on, verify the MetroCluster configuration from both sites.

To verify the configuration, repeat step 1.

8. Perform power cycle checks.

- a. Verify that all sync-source SVMs are online:

```
vserver show
```

- b. Start any sync-source SVMs that are not online:

```
vserver start
```

Maintenance procedures for MetroCluster IP configurations

IP switch maintenance and replacement

Replace an IP switch or change the use of existing MetroCluster IP switches

You might need to replace a failed switch, upgrade or downgrade a switch, or change the use of existing MetroCluster IP switches.

About this task

This procedure applies when you are using NetApp-validated switches. If you are using MetroCluster-compliant switches, refer to the switch vendor.

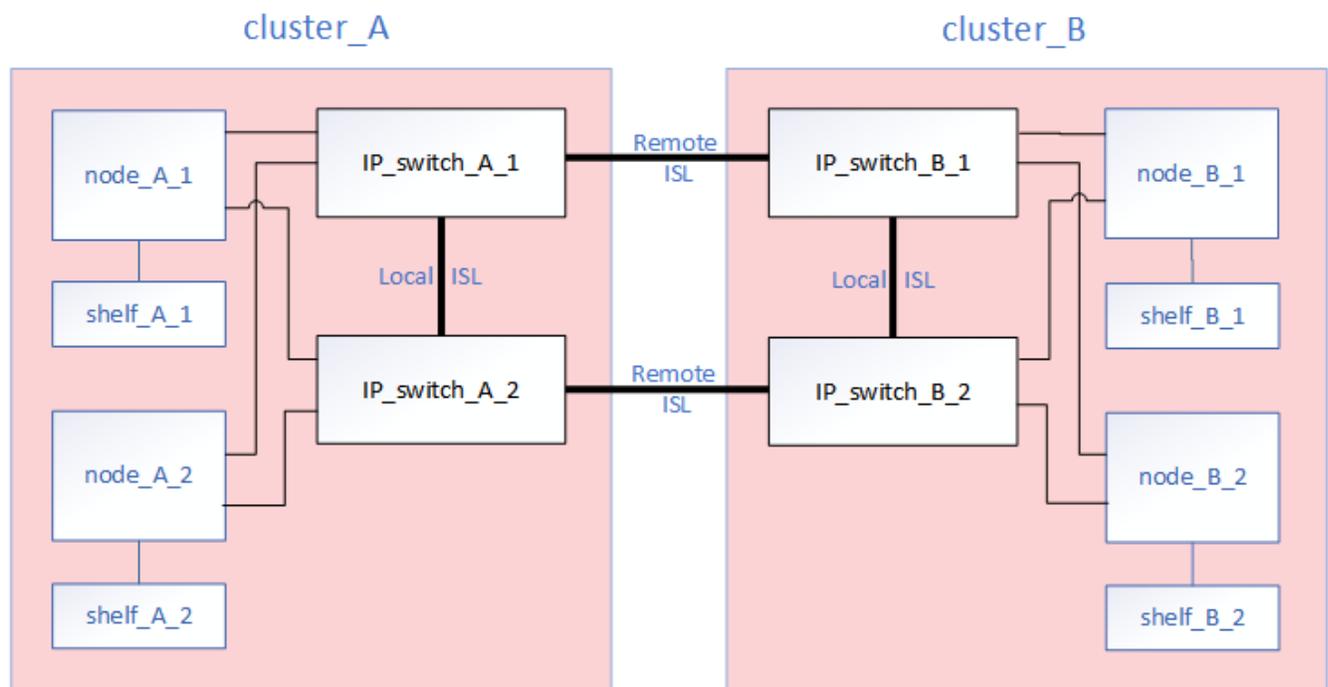
Enable console logging before performing this task.

This procedure supports the following conversions:

- Changing the switch vendor, type, or both. The new switch can be the same as the old switch when a switch has failed, or you can change the switch type (upgrade or downgrade the switch).

For example, to expand a MetroCluster IP configuration from a single four-node configuration using AFF A400 controllers and BES-53248 switches to an eight-node configuration using AFF A400 controllers, you must change the switches to a supported type for the configuration because BES-53248 switches are not supported in the new configuration.

If you want to replace a failed switch with the same type of switch, you only replace the failed switch. If you want to upgrade or downgrade a switch, you must adjust two switches that are in the same network. Two switches are in the same network when they are connected with an inter-switch link (ISL) and are not located at the same site. For example, Network 1 includes IP_switch_A_1 and IP_switch_B_1, and Network 2 includes IP_switch_A_2 and IP_switch_B_2, as shown in the diagram below:



If you replace a switch or upgrade to different switches, then you can pre-configure the switches by installing the switch firmware and RCF file.

- Convert a MetroCluster IP configuration to a MetroCluster IP configuration using shared storage MetroCluster switches.

For example, if you have a regular MetroCluster IP configuration using AFF A700 controllers and you want to reconfigure the MetroCluster to connect NS224 shelves to the same switches.



- If you are adding or removing shelves in a MetroCluster IP configuration using shared storage MetroCluster IP switches, follow the steps in [Adding shelves to a MetroCluster IP using shared storage MetroCluster switches](#)
- Your MetroCluster IP configuration might already directly connect to NS224 shelves or to dedicated storage switches.

Port usage worksheet

The following is an example worksheet for converting a MetroCluster IP configuration to a shared storage configuration connecting two NS224 shelves using the existing switches.

Worksheet definitions:

- Existing configuration: The cabling of the existing MetroCluster configuration.
- New configuration with NS224 shelves: The target configuration where the switches are shared between storage and the MetroCluster.

The highlighted fields in this worksheet indicate the following:

- Green: You do not need to change the cabling.
- Yellow: You must move ports with the same or a different configuration.
- Blue: Ports that are new connections.

PORT USAGE OVERVIEW

Example of expanding an existing 4Node MetroCluster with 2x NS224 shelves and changing the ISL's from 10G to 40/100G

Switch port	Existing configuration			New configuration with NS224 shelves		
	Port use	IP_switch_x_1	IP_switch_x_2	Port use	IP_switch_x_1	IP_switch_x_2
1	MetroCluster 1, Local Cluster Interface	Cluster Port 'A'	Cluster Port 'B'	MetroCluster 1, Local Cluster Interface	Cluster Port 'A'	Cluster Port 'B'
2		Cluster Port 'A'	Cluster Port 'B'		Cluster Port 'A'	Cluster Port 'B'
3						
4						
5				Storage shelf 1 (9)	NSM-A, e0a	NSM-A, e0b
6					NSM-B, e0a	NSM-B, e0b
7	ISL, Local Cluster native speed / 100G	ISL, Local Cluster		ISL, Local Cluster native speed / 100G	ISL, Local Cluster	
8						
9	MetroCluster 1, MetroCluster interface	Port 'A'	Port 'B'	MetroCluster 1, MetroCluster interface	Port 'A'	Port 'B'
10		Port 'A'	Port 'B'		Port 'A'	Port 'B'
11						
12						
13				ISL, MetroCluster, native speed 40G / 100G breakout mode 10G	Remote ISL, 2x 40/100G	Remote ISL, 2x 40/100G
14						
15						
16						
17				MetroCluster 1, Storage Interface	Storage Port 'A'	Storage Port 'B'
18					Storage Port 'A'	Storage Port 'B'
19						
20						
21	ISL, MetroCluster breakout mode 10G	Remote ISL, 10G	Remote ISL, 10G	Storage shelf 2 (8)	NSM-A, e0a	NSM-A, e0b
22					NSM-B, e0a	NSM-B, e0b
23						
24						
25						
26						
27						
28						
29						
30						
31						
32						
33						
34						
35						
36						

Steps

1. Check the health of the configuration.
 - a. Check that the MetroCluster is configured and in normal mode on each cluster: **metrocluster show**

```
cluster_A::> metrocluster show
Cluster                               Entry Name                               State
-----                               -
Local: cluster_A                      Configuration state configured
Mode                                   normal
AUSO Failure Domain auso-on-cluster-
disaster
Remote: cluster_B                     Configuration state configured
Mode                                   normal
AUSO Failure Domain auso-on-cluster-
disaster
```

- b. Check that mirroring is enabled on each node: **metrocluster node show**

```
cluster_A::> metrocluster node show
DR                                     Configuration DR
Group Cluster Node                    State           Mirroring Mode
-----
1   cluster_A
    node_A_1    configured     enabled   normal
    cluster_B
    node_B_1    configured     enabled   normal
2 entries were displayed.
```

- c. Check that the MetroCluster components are healthy: **metrocluster check run**

```
cluster_A::> metrocluster check run
```

```
Last Checked On: 10/1/2014 16:03:37
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok

4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results.

To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

d. Check that there are no health alerts: **system health alert show**

2. Configure the new switch before installation.

If you are reusing existing switches, go to [Step 4](#).



If you are upgrading or downgrading the switches, you must configure all the switches in the network.

Follow the steps in the section *Configuring the IP switches* in the [MetroCluster IP installation and configuration](#).

Make sure that you apply the correct RCF file for switch `_A_1`, `_A_2`, `_B_1` or `_B_2`. If the new switch is the same as the old switch, you need to apply the same RCF file.

If you upgrade or downgrade a switch, apply the latest supported RCF file for the new switch.

3. Run the port show command to view information about the network ports:

network port show

a. Modify all cluster LIFs to disable auto-revert:

```
network interface modify -vserver <vserver_name> -lif <lif_name>  
-auto-revert false
```

4. Disconnect the connections from the old switch.



You only disconnect connections that are not using the same port in the old and new configurations. If you are using new switches, you must disconnect all connections.

Remove the connections in the following order:

- a. Disconnect the local cluster interfaces
- b. Disconnect the local cluster ISLs
- c. Disconnect the MetroCluster IP interfaces
- d. Disconnect the MetroCluster ISLs

In the example [Port usage worksheet](#), the switches do not change. The MetroCluster ISLs are relocated and must be disconnected. You do not need to disconnect the connections marked in green on the worksheet.

5. If you are using new switches, power off the old switch, remove the cables, and physically remove the old switch.

If you are reusing existing switches, go to [Step 6](#).



Do **not** cable the new switches except for the management interface (if used).

6. Configure the existing switches.

If you have pre-configured the switches already, you can skip this step.

To configure the existing switches, follow the steps to install and upgrade the firmware and RCF files:

- [Upgrading firmware on MetroCluster IP switches](#)
- [Upgrade RCF files on MetroCluster IP switches](#)

7. Cable the switches.

You can follow the steps in the *Cabling the IP switches* section in [MetroCluster IP installation and configuration](#).

Cable the switches in the following order (if required):

- a. Cable the ISLs to the remote site.
- b. Cable the MetroCluster IP interfaces.
- c. Cable the local cluster interfaces.



- The used ports might be different from those on the old switch if the switch type is different.
If you are upgrading or downgrading the switches, do **NOT** cable the local ISLs. Only cable the local ISLs if you are upgrading or downgrading the switches in the second network and both switches at one site are the same type and cabling.
- If you are upgrading Switch-A1 and Switch-B1, you must perform steps 1 to 6 for switches Switch-A2 and Switch-B2.

8. Finalize the local cluster cabling.

- a. If the local cluster interfaces are connected to a switch:
 - i. Cable the local cluster ISLs.
 - b. If the local cluster interfaces are **not** connected to a switch:
 - i. Use the [Migrate to a switched NetApp cluster environment](#) procedure to convert a switchless cluster to a switched cluster. Use the ports indicated in [MetroCluster IP installation and configuration](#) or the RCF cabling files to connect the local cluster interface.
9. Power up the switch or switches.

If the new switch is the same, power up the new switch. If you are upgrading or downgrading the switches, then power up both switches. The configuration can operate with two different switches at each site until the second network is updated.

10. Verify that the MetroCluster configuration is healthy by repeating [Step 1](#).

If you are upgrading or downgrading the switches in the first network, you might see some alerts related to local clustering.



If you upgrade or downgrade the networks, then repeat all of the steps for the second network.

11. Modify all cluster LIFs to re-enable auto-revert:

```
network interface modify -vserver <vserver_name> -lif <lif_name> -auto
-revert true
```

12. Revert all cluster LIFs that are not currently on their home ports to their home ports:

```
network interface revert -vserver * -lif *
```

13. Optionally, move the NS224 shelves.

If you are reconfiguring a MetroCluster IP configuration that does not connect NS224 shelves to the MetroCluster IP switches, use the appropriate procedure to add or move the NS224 shelves:

- [Adding shelves to a MetroCluster IP using shared storage MetroCluster switches](#)
- [Migrate from a switchless cluster with direct-attached storage](#)
- [Migrate from a switchless configuration with switch-attached storage by reusing the storage switches](#)

Online or offline MetroCluster IP interface ports

When you perform maintenance tasks, you might need to bring a MetroCluster IP interface port offline or online.

About this task

[Enable console logging](#) before performing this task.

Steps

You can use the following steps to bring a MetroCluster IP interface port online or take it offline.

1. Set the privilege level to advanced.

```
set -privilege advanced
```

Example output

```
Cluster A_1::> set -privilege advanced
Warning: These advanced commands are potentially dangerous; use them
only when
        directed to do so by NetApp personnel.
Do you want to continue? {y|n}: y
```

2. Take the MetroCluster IP interface port offline.

```
system ha interconnect link off -node <node_name> -link <link_num, 0 or
1>
```

Example output

```
Cluster_A1::*> system ha interconnect link off -node node-a1 -link 0
```

- a. Verify the MetroCluster IP interface is offline.

```
Cluster_A1::*> system ha interconnect port show
```

Example output

```

Cluster_A1::*> system ha interconnect port show

```

Physical Node Down	Link Active Link	Monitor	Port	Physical Layer State	Link Layer State	Physical Link Up	Link
node-a1	off			0 disabled	down	4	
3 false				1 linkup	active	4	
2 true				0 linkup	active	4	
node-a2	off			1 linkup	active	4	
2 true							

2 entries were displayed.

3. Bring the MetroCluster IP interface port online.

```

system ha interconnect link on -node <node_name> -link <link_num, 0 or 1>

```

Example output

```

Cluster_A1::*> system ha interconnect link on -node node-a1 -link 0

```

a. Verify the MetroCluster IP interface port is online.

```

Cluster_A1::*> system ha interconnect port show

```

Example output

```

Cluster_A1::*> system ha interconnect port show
                Physical Link
                Layer   Layer   Physical
Physical Active
Node           Monitor Port  State  State  Link Up  Link
Down  Link
-----
node-a1        off
                0  linkup  active  5
3  true
                1  linkup  active  4
2  true
node-a2        off
                0  linkup  active  4
2  true
                1  linkup  active  4
2  true
2 entries were displayed.

```

Upgrade firmware on MetroCluster IP switches

You might need to upgrade the firmware on a MetroCluster IP switch.

Verify that the RCF is supported

When you change ONTAP version or the switch firmware version, you should verify that you have a reference configuration file (RCF) that is supported for that version. If you use the [RcfFileGenerator](#) tool, the correct RCF is generated for your configuration.

Steps

1. Use the following commands from the switches to verify the version of the RCF:

From this switch...	Issue this command...
Broadcom switch	(IP_switch_A_1) # show clibanner
Cisco switch	IP_switch_A_1# show banner motd
NVIDIA SN2100 switch	cumulus@mcc1:mgmt:~\$ nv config find message

Locate the line in the command output that indicates the RCF version. For example, the following output from a Cisco switch indicates that the RCF version is “v1.80”.

```
Filename : NX3232_v1.80_Switch-A2.txt
```

2. To check which files are supported for a specific ONTAP version, switch, and platform, use the [RcfFileGenerator for MetroCluster IP](#). If you can generate the RCF for the configuration that you have or that you want to upgrade to, then it is supported.
3. To verify that the switch firmware is supported, refer to the following:
 - [Hardware Universe](#)
 - [NetApp Interoperability Matrix](#)

Upgrade the switch firmware

About this task

You must repeat this task on each of the switches in succession.

[Enable console logging](#) before performing this task.

Steps

1. Check the health of the configuration.
 - a. Check that the MetroCluster is configured and in normal mode on each cluster:

```
metrocluster show
```

```
cluster_A::> metrocluster show
Cluster                Entry Name                State
-----
Local: cluster_A      Configuration state    configured
                      Mode                    normal
                      AUSO Failure Domain   auso-on-cluster-
disaster
Remote: cluster_B     Configuration state    configured
                      Mode                    normal
                      AUSO Failure Domain   auso-on-cluster-
disaster
```

- b. Check that mirroring is enabled on each node:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR                               Configuration  DR
Group Cluster Node              State          Mirroring Mode
-----
1      cluster_A
      node_A_1      configured    enabled      normal
      cluster_B
      node_B_1      configured    enabled      normal
2 entries were displayed.

```

c. Check that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

d. After the metrocluster check run operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```

cluster_A::*> metrocluster check show

Component          Result
-----
nodes              ok
lifs               ok
config-replication ok
aggregates        ok
clusters           ok
connections        ok
volumes           ok
7 entries were displayed.

```

e. Check that there are no health alerts:

```
system health alert show
```

2. Install the software on the first switch.



You must install the switch software on the switches in the following order: switch_A_1, switch_B_1, switch_A_2, switch_B_2.

Follow the steps for installing switch software in the relevant topic depending on whether the switch type is Broadcom, Cisco, or NVIDIA:

- [Download and install the Broadcom switch EFOS software](#)
- [Download and install the Cisco switch NX-OS software](#)
- [Download and install the NVIDIA SN2100 switch Cumulus software](#)

3. Repeat the previous step for each of the switches.
4. Repeat [Step 1](#) to check the health of the configuration.

Upgrade RCF files on MetroCluster IP switches

You might need to upgrade a reference configuration file (RCF) file on a MetroCluster IP switch. For example, if the RCF version that you are running on the switches is not supported by the ONTAP version, the switch firmware version, or both.

Before you begin

- If you are installing new switch firmware, you must install the switch firmware before upgrading the RCF file.
- Before you upgrade the RCF, [verify that the RCF is supported](#).
- [Enable console logging](#) before performing this task.

About this task

- This procedure disrupts traffic on the switch where the RCF file is upgraded. Traffic resumes when the new RCF file is applied.
- Perform the steps on one switch at a time, in the following order: Switch_A_1, Switch_B_1, Switch_A_2, Switch_B_2.

Steps

1. Verify the health of the configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- a. After the `metrocluster check run` operation completes, run `metrocluster check show` to view the results.

After approximately five minutes, the following results are displayed:

```

-----
::*> metrocluster check show

Component          Result
-----
nodes              ok
lifs               ok
config-replication ok
aggregates        ok
clusters           ok
connections        ok
volumes           ok
7 entries were displayed.

```

b. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id 38
```

c. Verify that there are no health alerts:

```
system health alert show
```

2. Prepare the IP switches for the application of the new RCF files.

Follow the steps for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

3. Download and install the IP RCF file, depending on your switch vendor.

- [Download and install the Broadcom IP RCF files](#)
- [Download and install the Cisco IP RCF files](#)
- [Download and install the NVIDIA IP RCF files](#)



If you have an L2 shared or L3 network configuration, you might need to adjust the ISL ports on the intermediate/customer switches. The switchport mode might change from 'access' to 'trunk' mode. Only proceed to upgrade the second switch pair (A_2, B_2) if the network connectivity between switches A_1 and B_1 is fully operational and the network is healthy.

Upgrade RCF files on Cisco IP switches using CleanUpFiles

You might need to upgrade an RCF file on a Cisco IP switch. For example, an ONTAP upgrade or a switch firmware upgrade both require a new RCF file.

About this task

- Beginning with RcfFileGenerator version 1.4a, there is a new option to change (upgrade, downgrade, or replace) the switch configuration on Cisco IP switches without the need to perform a 'write erase'.
- [Enable console logging](#) before performing this task.
- The Cisco 9336C-FX2 switch has two different switch storage types that are named differently in the RCF. Use the following table to determine the correct Cisco 9336C-FX2 storage type for your configuration:

If you are connecting the following storage...	Choose the Cisco 9336C-FX2 storage type...	Sample RCF file banner/MOTD
<ul style="list-style-type: none"> • Directly connected SAS shelves • Directly connected NVMe shelves • NVMe shelves connected to dedicated storage switches 	9336C-FX2 – Direct Storage only	* Switch : NX9336C (direct storage, L2 Networks, direct ISL)
<ul style="list-style-type: none"> • Directly connected SAS shelves • NVMe shelves connected to the MetroCluster IP switches <div style="border-left: 1px solid #ccc; padding-left: 10px; margin-top: 10px;">  At least one Ethernet connected NVMe shelf is required </div>	9336C-FX2 – SAS and Ethernet storage	* Switch : NX9336C (SAS and Ethernet storage, L2 Networks, direct ISL)

Before you begin

You can use this method if your configuration meets the following requirements:

- The standard RCF configuration is applied.
- The [RcfFileGenerator](#) must be able to create the same RCF file that is applied, with the same version and configuration (platforms, VLANs).
- The RCF file that is applied was not provided by NetApp for a special configuration.
- The RCF file was not altered before it was applied.
- The steps to reset the switch to factory defaults were followed before applying the current RCF file.
- No changes were made to the switch(port) configuration after the RCF was applied.

If you do not meet these requirements, then you cannot use the CleanUpFiles that are created when generating the RCF files. However, you can leverage the function to create generic CleanUpFiles — the cleanup using this method is derived from the output of `show running-config` and is best practice.



You must update the switches in the following order: Switch_A_1, Switch_B_1, Switch_A_2, Switch_B_2. Or, you can update the switches Switch_A_1 and Switch_B_1 at the same time followed by switches Switch_A_2 and Switch_B_2.

Steps

1. Determine the current RCF file version, and which ports and VLANs are used: `IP_switch_A_1# show banner motd`



You need to get this information from all four switches and complete the following information table.

```
* NetApp Reference Configuration File (RCF)
*
* Switch : NX9336C (SAS storage, L2 Networks, direct ISL)
* Filename : NX9336_v1.81_Switch-A1.txt
* Date : Generator version: v1.3c_2022-02-24_001, file creation time:
2021-05-11, 18:20:50
*
* Platforms : MetroCluster 1 : FAS8300, AFF-A400, FAS8700
*             MetroCluster 2 : AFF-A320, FAS9000, AFF-A700, AFF-A800
* Port Usage:
* Ports 1- 2: Intra-Cluster Node Ports, Cluster: MetroCluster 1, VLAN
111
* Ports 3- 4: Intra-Cluster Node Ports, Cluster: MetroCluster 2, VLAN
151
* Ports 5- 6: Ports not used
* Ports 7- 8: Intra-Cluster ISL Ports, local cluster, VLAN 111, 151
* Ports 9-10: MetroCluster 1, Node Ports, VLAN 119
* Ports 11-12: MetroCluster 2, Node Ports, VLAN 159
* Ports 13-14: Ports not used
* Ports 15-20: MetroCluster-IP ISL Ports, VLAN 119, 159, Port Channel 10
* Ports 21-24: MetroCluster-IP ISL Ports, VLAN 119, 159, Port Channel
11, breakout mode 10gx4
* Ports 25-30: Ports not used
* Ports 31-36: Ports not used
*
#
IP_switch_A_1#
```

From this output, you must collect the information shown in the following two tables.

Generic information	MetroCluster	Data
RCF file version		1.81
Switch type		NX9336
Network typology		L2 Networks, direct ISL
Storage type		SAS storage

Platforms	1	AFF A400
	2	FAS9000

VLAN information	Network	MetroCluster configuration	Switchports	Site A	Site B
VLAN local cluster	Network 1	1	1, 2	111	222
		2	3, 4	151	251
	Network 2	1	1, 2	111	222
		2	3, 4	151	251
VLAN MetroCluster	Network 1	1	9, 10	119	119
		2	11, 12	159	159
	Network 2	1	9, 10	219	219
		2	11, 12	259	259

2. Create the RCF files and CleanUpFiles, or create generic CleanUpFiles for the current configuration.

If your configuration meets the requirements outlined in the prerequisites, select **Option 1**.

If your configuration does **not** meet the requirements outlined in the prerequisites, select **Option 2**.

Option 1: Create the RCF files and CleanUpFiles

Use this procedure if the configuration meets the requirements.

Steps

- a. Use the RcfFileGenerator 1.4a (or later) to create the RCF files with the information that you retrieved in Step 1. The new version of the RcfFileGenerator creates an additional set of CleanUpFiles that you can use to revert some configuration and prepare the switch to apply a new RCF configuration.
- b. Compare the banner motd with the RCF files that are currently applied. The platform types, switch type, port and VLAN usage must be the same.



You must use the CleanUpFiles from the same version as the RCF file and for the exact same configuration. Using any CleanUpFile will not work and might require a full reset of the switch.



The ONTAP version the RCF file is created for is not relevant. Only the RCF file version is important.



The RCF file (even it is the same version) might list fewer or more platforms. Make sure that your platform is listed.

Option 2: Create generic CleanUpFiles

Use this procedure if the configuration does **not** meet all the requirements.

Steps

- a. Retrieve the output of `show running-config` from each switch.
- b. Open the RcfFileGenerator tool and click 'Create generic CleanUpFiles' at the bottom of the window
- c. Copy the output that you retrieved in Step 1 from 'one' switch into the upper window. You can remove or leave the default output.
- d. Click 'Create CUF files'.
- e. Copy the output from the lower window into a text file (this file is the CleanUpFile).
- f. Repeat Steps c, d, and e for all switches in the configuration.

At the end of this procedure, you should have four text files, one for each switch. You can use these files in the same way as the CleanUpFiles that you can create by using Option 1.

3. Create the 'new' RCF files for the new configuration.
Create these files in the same way that you created the files in the previous step, except choose the respective ONTAP and RCF file version.

After completing this step you should have two sets of RCF files, each set consisting of twelve files.

4. Download the files to the bootflash.
 - a. Download the CleanUpFiles that you created in [Create the RCF files and CleanUpFiles, or create generic CleanUpFiles for the current configuration](#)



This CleanUpFile is for the current RCF file that is applied and **NOT** for the new RCF that you want to upgrade to.

Example CleanUpFile for Switch-A1: Cleanup_NX9336_v1.81_Switch-A1.txt

- b. Download the 'new' RCF files that you created in [Create the 'new' RCF files for the new configuration](#).

Example RCF file for Switch-A1: NX9336_v1.90_Switch-A1.txt

- c. Download the CleanUpFiles that you created in [Create the 'new' RCF files for the new configuration](#). This step is optional — you can use the file in future to update the switch configuration. It matches the currently applied configuration.

Example CleanUpFile for Switch-A1: Cleanup_NX9336_v1.90_Switch-A1.txt



You must use the CleanUpFile for the correct (matching) RCF version. If you use a CleanUpFile for a different RCF version, or a different configuration then the cleanup of the configuration might not work correctly.

The following example copies the three files to the bootflash:

```
IP_switch_A_1# copy sftp://user@50.50.50.50/RcfFiles/NX9336-direct-
SAS_v1.81_MetroCluster-
IP_L2Direct_A400FAS8700_xxx_xxx_xxx_xxx/Cleanup_NX9336_v1.81_Switch-
A1.txt bootflash:
IP_switch_A_1# copy sftp://user@50.50.50.50/RcfFiles/NX9336-direct-
SAS_v1.90_MetroCluster-
IP_L2Direct_A400FAS8700A900FAS9500_xxx_xxx_xxx_xxxNX9336_v1.90//NX933
6_v1.90_Switch-A1.txt bootflash:
IP_switch_A_1# copy sftp://user@50.50.50.50/RcfFiles/NX9336-direct-
SAS_v1.90_MetroCluster-
IP_L2Direct_A400FAS8700A900FAS9500_xxx_xxx_xxx_xxxNX9336_v1.90//Clean
up_NX9336_v1.90_Switch-A1.txt bootflash:
```



You are prompted to specify Virtual Routing and Forwarding (VRF).

5. Apply the CleanUpFile or generic CleanUpFile.

Some of the configuration is reverted and switchports go 'offline'.

- a. Confirm that there are no pending changes to the startup configuration: `show running-config diff`

```
IP_switch_A_1# show running-config diff
IP_switch_A_1#
```

6. If you see system output, save the running configuration to the startup configuration: `copy running-`

```
config startup-config
```



System output indicates that the startup configuration and running configuration are different and pending changes. If you do not save the pending changes, you are unable to roll back using a reload of the switch.

a. Apply the CleanUpFile:

```
IP_switch_A_1# copy bootflash:Cleanup_NX9336_v1.81_Switch-A1.txt
running-config

IP_switch_A_1#
```



The script might take a while to return to the switch prompt. No output is expected.

7. View the running configuration to verify that the configuration is cleared: `show running-config`

The current configuration should show:

- No class maps and IP access lists are configured
- No policy maps are configured
- No service policies are configured
- No port-profiles are configured
- All Ethernet interfaces (except mgmt0 which should not show any configuration, and only VLAN 1 should be configured).

If you find that any of the above items are configured, you might not be able to apply a new RCF file configuration. However, you can revert to the previous configuration by reloading the switch **without** saving the running configuration to the startup configuration. The switch will come up with the previous configuration.

8. Apply the RCF file and verify that the ports are online.

a. Apply the RCF files.

```
IP_switch_A_1# copy bootflash:NX9336_v1.90-X2_Switch-A1.txt running-
config
```



Some warning messages appear while applying the configuration. Error messages are generally not expected. However, if you are logged in using SSH, you might receive the following error: `Error: Can't disable/re-enable ssh:Current user is logged in through ssh`

b. After the configuration is applied, verify that the cluster and MetroCluster ports are coming online with one of the following commands, `show interface brief`, `show cdp neighbors`, or `show lldp neighbors`



If you changed the VLAN for the local cluster and you upgraded the first switch at the site, then cluster health monitoring might not report the state as 'healthy' because the VLANs from the old and new configurations do not match. After the second switch is updated, the state should return to healthy.

If the configuration is not applied correctly, or you do not want to keep the configuration, you can revert to the previous configuration by reloading the switch **without** saving the running configuration to startup configuration. The switch will come up with the previous configuration.

9. Save the configuration and reload the switch.

```
IP_switch_A_1# copy running-config startup-config
```

```
IP_switch_A_1# reload
```

Renaming a Cisco IP switch

You might need to rename a Cisco IP switch to provide consistent naming throughout your configuration.

About this task

- In the examples in this task, the switch name is changed from `myswitch` to `IP_switch_A_1`.
- [Enable console logging](#) before performing this task.

Steps

1. Enter global configuration mode:

```
configure terminal
```

The following example shows the configuration mode prompt. Both prompts show the switch name of `myswitch`.

```
myswitch# configure terminal  
myswitch(config)#
```

2. Rename the switch:

```
switchname new-switch-name
```

If you are renaming both switches in the network, use the same command on each switch.

The CLI prompt changes to reflect the new name:

```
myswitch(config)# switchname IP_switch_A_1  
IP_switch_A_1(config)#
```

3. Exit configuration mode:

exit

The top-level switch prompt is displayed:

```
IP_switch_A_1(config)# exit
IP_switch_A_1#
```

4. Copy the current running configuration to the startup configuration file:

copy running-config startup-config

5. Verify that the switch name change is visible from the ONTAP cluster prompt.

Note that the new switch name is shown, and the old switch name (`myswitch`) does not appear.

- a. Enter advanced privilege mode, pressing **y** when prompted:

set -privilege advanced

- b. Display the attached devices:

network device-discovery show

- c. Return to admin privilege mode:

set -privilege admin

The following example shows that the switch appears with the new name, `IP_switch_A_1`:

```
cluster_A::storage show> set advanced
```

Warning: These advanced commands are potentially dangerous; use them only when directed to do so by NetApp personnel.

```
Do you want to continue? {y|n}: y
```

```
cluster_A::storage show*> network device-discovery show
```

Node/	Local	Discovered	Interface	
Protocol	Port	Device		
Platform				

node_A_2/cdp				
	e0M	LF01-410J53.mycompany.com (SAL18516DZY)	Ethernet125/1/28	N9K-
C9372PX				
	e1a	IP_switch_A_1 (FOC21211RBU)	Ethernet1/2	N3K-
C3232C				
	e1b	IP_switch_A_1 (FOC21211RBU)	Ethernet1/10	N3K-
C3232C				
.				
.			Ethernet1/18	N9K-
C9372PX				
node_A_1/cdp				
	e0M	LF01-410J53.mycompany.com (SAL18516DZY)	Ethernet125/1/26	N9K-
C9372PX				
	e0a	IP_switch_A_2 (FOC21211RB5)	Ethernet1/1	N3K-
C3232C				
	e0b	IP_switch_A_2 (FOC21211RB5)	Ethernet1/9	N3K-
C3232C				
	e1a	IP_switch_A_1 (FOC21211RBU)		
.				
.				
.				

16 entries were displayed.

Add, remove, or change ISL ports nondisruptively on Cisco IP switches

You might need to add, remove, or change ISL ports on Cisco IP switches. You can

convert dedicated ISL ports to shared ISL ports, or change the speed of ISL ports on a Cisco IP switch.

About this task

If you are converting dedicated ISL ports to shared ISL ports, ensure the new ports meet the [Requirements for shared ISL ports](#).

You must complete all the steps on both switches to ensure ISL connectivity.

The following procedure assumes you are replacing a 10-Gb ISL connected at switch port Eth1/24/1 with two 100-Gb ISLs that are connected to switch ports 17 and 18.



If you are using a Cisco 9336C-FX2 switch in a shared configuration connecting NS224 shelves, changing the ISLs might require a new RCF file. You do not require a new RCF file if your current and new ISL speed is 40Gbps and 100Gbps. All other changes to ISL speed requires a new RCF file. For example, changing the ISL speed from 40Gbps to 100Gbps does not require a new RCF file, but changing the ISL speed from 10Gbps to 40Gbps requires a new RCF file.

Before you begin

Refer to the **Switches** section of the [NetApp Hardware Universe](#) to verify the supported transceivers.

[Enable console logging](#) before performing this task.

Steps

1. Disable the ISL ports of the ISLs on both switches in the fabric that you want to change.



You only need to disable the current ISL ports if you are moving them to a different port, or the speed of the ISL is changing. If you are adding an ISL port with the same speed as the existing ISLs, go to Step 3.

You must enter only one configuration command for each line and press Ctrl-Z after you have entered all the commands, as shown in the following example:

```
switch_A_1# conf t
switch_A_1(config)# int eth1/24/1
switch_A_1(config-if)# shut
switch_A_1(config-if)#
switch_A_1#

switch_B_1# conf t
switch_B_1(config)# int eth1/24/1
switch_B_1(config-if)# shut
switch_B_1(config-if)#
switch_B_1#
```

2. Remove the existing cables and transceivers.
3. Change the ISL port as required.



If you are using Cisco 9336C-FX2 switches in a shared configuration connecting NS224 shelves, and you need to upgrade the RCF file and apply the new configuration for the new ISL ports, follow the steps to [upgrade the RCF files on MetroCluster IP switches](#).

Option	Step
To change the speed of an ISL port...	Cable the new ISLs to the designated ports according to their speeds. You must ensure that these ISL ports for your switch are listed in the <i>MetroCluster IP Installation and Configuration</i> .
To add an ISL...	Insert QFSPs into the ports you are adding as ISL ports. Ensure they are listed in the <i>MetroCluster IP Installation and Configuration</i> and cable them accordingly.

4. Enable all ISL ports (if not enabled) on both switches in the fabric beginning with the following command:

```
switch_A_1# conf t
```

You must enter only one configuration command per line and press Ctrl-Z after you have entered all the commands:

```
switch_A_1# conf t
switch_A_1(config)# int eth1/17
switch_A_1(config-if)# no shut
switch_A_1(config-if)# int eth1/18
switch_A_1(config-if)# no shut
switch_A_1(config-if)#
switch_A_1#
switch_A_1# copy running-config startup-config

switch_B_1# conf t
switch_B_1(config)# int eth1/17
switch_B_1(config-if)# no shut
switch_B_1(config-if)# int eth1/18
switch_B_1(config-if)# no shut
switch_B_1(config-if)#
switch_B_1#
switch_B_1# copy running-config startup-config
```

5. Verify that the ISLs and port channels for the ISLs are established between both switches:

```
switch_A_1# show int brief
```

You should see the ISL interfaces in the command output as shown in the following example:

```

Switch_A_1# show interface brief
-----
-----
Ethernet          VLAN    Type Mode   Status Reason          Speed
Port
Interface
Ch #
-----
-----
Eth1/17           1       eth  access down   XCVR not inserted
auto(D) --
Eth1/18           1       eth  access down   XCVR not inserted
auto(D) --
-----
-----
Port-channel      VLAN    Type Mode   Status Reason
Speed  Protocol
Interface
-----
-----
Po10              1       eth  trunk  up     none
a-100G(D) lacp
Po11              1       eth  trunk  up     none
a-100G(D) lacp

```

6. Repeat the procedure for fabric 2.

Identifying storage in a MetroCluster IP configuration

If you need to replace a drive or shelf module, you first need to identify the location.

Identification of local and remote shelves

When you view shelf information from a MetroCluster site, all remote drives are on 0m, the virtual iSCSI host adapter. This means that the drives are accessed via the MetroCluster IP interfaces. All other drives are local.

After identifying whether a shelf is remote (on 0m), you can further identify the drive or shelf by the serial number or, depending on shelf ID assignments in your configuration, by shelf ID.



In MetroCluster IP configurations running ONTAP 9.4, the shelf ID is not required to be unique between the MetroCluster sites. This includes both internal shelves (0) and external shelves. The serial number is consistent when viewed from any node on either MetroCluster site.

Shelf IDs should be unique within the disaster recovery (DR) group except for the internal shelf.

With the drive or shelf module identified, you can replace the component using the appropriate procedure.

Example of sysconfig -a output

The following example uses the `sysconfig -a` command to show the devices on a node in the MetroCluster IP configuration. This node has the following shelves and devices attached:

- slot 0: Internal drives (local drives)
- slot 3: External shelf ID 75 and 76 (local drives)
- slot 0: Virtual iSCSI host adapter 0m (remote drives)

```
node_A_1> run local sysconfig -a

NetApp Release R9.4:  Sun Mar 18 04:14:58 PDT 2018
System ID: 1111111111 (node_A_1); partner ID: 2222222222 (node_A_2)
System Serial Number: serial-number (node_A_1)
.
.
.
slot 0: NVMe Disks
           0      : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500528)
           1      : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500735)
           2      : NETAPP  X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J501165)
.
.
.
slot 3: SAS Host Adapter 3a (PMC-Sierra PM8072 rev. C, SAS, <UP>)
MFG Part Number:  Microsemi Corp. 110-03801 rev. A0
Part number:      111-03801+A0
Serial number:    7A1063AF14B
Date Code:        20170320
Firmware rev:    03.08.09.00
Base WWN:         5:0000d1:702e69e:80
Phy State:        [12] Enabled, 12.0 Gb/s
                  [13] Enabled, 12.0 Gb/s
                  [14] Enabled, 12.0 Gb/s
                  [15] Enabled, 12.0 Gb/s
Mini-SAS HD Vendor:  Molex Inc.
Mini-SAS HD Part Number:  112-00436+A0
Mini-SAS HD Type:        Passive Copper (unequalized) 0.5m ID:00
Mini-SAS HD Serial Number: 614130640
                        75.0 : NETAPP  X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501805)
```

75.1 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502050)
75.2 : NETAPP X438_PHM2400MCTO NA04 381.3GB 520B/sect
(25M0A03WT2KA)
75.3 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501793)
75.4 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502158)

.
. .
.

Shelf 75: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220
Shelf 76: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220

slot 3: SAS Host Adapter 3c (PMC-Sierra PM8072 rev. C, SAS, <UP>)

MFG Part Number: Microsemi Corp. 110-03801 rev. A0

Part number: 111-03801+A0

Serial number: 7A1063AF14B

Date Code: 20170320

Firmware rev: 03.08.09.00

Base WWN: 5:0000d1:702e69e:88

Phy State: [0] Enabled, 12.0 Gb/s

[1] Enabled, 12.0 Gb/s

[2] Enabled, 12.0 Gb/s

[3] Enabled, 12.0 Gb/s

Mini-SAS HD Vendor: Molex Inc.

Mini-SAS HD Part Number: 112-00436+A0

Mini-SAS HD Type: Passive Copper (unequalized) 0.5m ID:00

Mini-SAS HD Serial Number: 614130691

75.0 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501805)

75.1 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502050)

75.2 : NETAPP X438_PHM2400MCTO NA04 381.3GB 520B/sect
(25M0A03WT2KA)

75.3 : NETAPP X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501793)

.
. .
.

Shelf 75: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220

Shelf 76: DS224-12 Firmware rev. IOM12 A: 0220 IOM12 B: 0220

slot 3: SAS Host Adapter 3d (PMC-Sierra PM8072 rev. C, SAS, <UP>)

MFG Part Number: Microsemi Corp. 110-03801 rev. A0

```

Part number:      111-03801+A0
Serial number:    7A1063AF14B
Date Code:       20170320
Firmware rev:    03.08.09.00
Base WWN:        5:0000d1:702e69e:8c
Phy State:       [4] Enabled, 12.0 Gb/s
                  [5] Enabled, 12.0 Gb/s
                  [6] Enabled, 12.0 Gb/s
                  [7] Enabled, 12.0 Gb/s

Mini-SAS HD Vendor:      Molex Inc.
Mini-SAS HD Part Number: 112-00436+A0
Mini-SAS HD Type:        Passive Copper (unequalized) 0.5m ID:01
Mini-SAS HD Serial Number: 614130690
                        75.0 : NETAPP    X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG501805)
                        75.1 : NETAPP    X438_S1633400AMD NA04 381.3GB 520B/sect
(S20KNYAG502050)
                        75.2 : NETAPP    X438_PHM2400MCTO NA04 381.3GB 520B/sect
(25M0A03WT2KA)
.
.
.
Shelf 75: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220
Shelf 76: DS224-12  Firmware rev. IOM12 A: 0220  IOM12 B: 0220

slot 4: Quad 10 Gigabit Ethernet Controller X710 SFP+
.
.
.
slot 0: Virtual iSCSI Host Adapter 0m
        0.0 : NETAPP    X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500690)
        0.1 : NETAPP    X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500571)
        0.2 : NETAPP    X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500323)
        0.3 : NETAPP    X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500724)
        0.4 : NETAPP    X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500734)
        0.5 : NETAPP    X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500598)
        0.12 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J501094)
        0.13 : NETAPP   X4001S172A1T9NTE NA01 1831.1GB 4160B/sect
(S3NBNX0J500519)

```

```
.  
. .  
Shelf 0: FS4483PSM3E  Firmware rev. PSM3E A: 0103  PSM3E B: 0103  
Shelf 35: DS224-12   Firmware rev. IOM12 A: 0220  IOM12 B: 0220  
Shelf 36: DS224-12   Firmware rev. IOM12 A: 0220  IOM12 B: 0220  
  
node_A_1::>
```

Adding shelves to a MetroCluster IP using shared Storage MetroCluster switches

You might need to add NS224 shelves to a MetroCluster using shared Storage MetroCluster switches.

Starting from ONTAP 9.10.1, you can add NS224 shelves from a MetroCluster using the shared Storage / MetroCluster switches. You can add more than one shelf at a time.

Before you begin

- Nodes must be running ONTAP 9.9.1 or later.
- All currently connected NS224 shelves must be attached to the same switches as the MetroCluster (shared Storage / MetroCluster switch configuration).
- This procedure cannot be used to convert a configuration with directly connected NS224 shelves or NS224 shelves attached to dedicated Ethernet switches to a configuration using shared Storage / MetroCluster switches.
- [Enable console logging](#) before performing this task.

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. To prevent automatic support case generation, send an Autosupport message to indicate the upgrade is underway.
 - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message "Maint=10h Adding  
or Removing NS224 shelves" _
```

This example specifies a 10 hour maintenance window. You might want to allow additional time, depending on your plan.

If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

Verifying the health of the MetroCluster configuration

You must verify the health and connectivity of the MetroCluster configuration prior to performing the transition.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- g. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

2. Verify that the cluster is healthy:

```
cluster show -vserver Cluster
```

```
cluster_A::> cluster show -vserver Cluster
Node           Health  Eligibility  Epsilon
-----
node_A_1       true   true         false
node_A_2       true   true         false

cluster_A::>
```

3. Verify that all cluster ports are up:

```
network port show -ipSPACE cluster
```

```
cluster_A::> network port show -ipSPACE cluster

Node: node_A_1-old

Port          IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
-----
e0a           Cluster      Cluster          up  9000    auto/10000 healthy
e0b           Cluster      Cluster          up  9000    auto/10000 healthy

Node: node_A_2-old

Port          IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
-----
e0a           Cluster      Cluster          up  9000    auto/10000 healthy
e0b           Cluster      Cluster          up  9000    auto/10000 healthy

4 entries were displayed.

cluster_A::>
```

4. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

```
cluster_A::> network interface show -vserver cluster
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
Cluster	node_A_1-old_clus1	up/up	169.254.209.69/16	node_A_1	e0a
true	node_A_1-old_clus2	up/up	169.254.49.125/16	node_A_1	e0b
true	node_A_2-old_clus1	up/up	169.254.47.194/16	node_A_2	e0a
true	node_A_2-old_clus2	up/up	169.254.19.183/16	node_A_2	e0b
true					

```
4 entries were displayed.
```

```
cluster_A::>
```

5. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

```

cluster_A::> network interface show -vserver Cluster -fields auto-revert

          Logical
Vserver  Interface      Auto-revert
-----  -
Cluster
          node_A_1-old_clus1
                        true
          node_A_1-old_clus2
                        true
          node_A_2-old_clus1
                        true
          node_A_2-old_clus2
                        true

          4 entries were displayed.

cluster_A::>

```

Applying the new RCF file to the switches



If your switch is already correctly configured, you can skip these next sections and go directly to [Configuring MACsec encryption on Cisco 9336C switches](#), if applicable or to [Connecting the new NS224 shelf](#).

- You must change the switch configuration to add shelves.
- You should review the cabling details at [Platform port assignments](#).
- You must use the **RcfFileGenerator** tool to create the RCF file for your configuration. The [RcfFileGenerator](#) also provides a per-port cabling overview for each switch. Make sure that you choose the correct number of shelves. There are additional files created along with the RCF file that provide a detailed cabling layout matching your specific options. Use this cabling overview to verify your cabling when cabling the new shelves.

Upgrading RCF files on MetroCluster IP switches

If you are installing new switch firmware, you must install the switch firmware before upgrading the RCF file.

This procedure disrupts traffic on the switch where the RCF file is upgraded. Traffic will resume once the new RCF file is applied.

Steps

1. Verify the health of the configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- a. After the `metrocluster check run` operation completes, run `metrocluster check show` to view the results.

After approximately five minutes, the following results are displayed:

```
-----  
::*> metrocluster check show  
  
Component          Result  
-----  
nodes              ok  
lifs               ok  
config-replication ok  
aggregates        ok  
clusters          ok  
connections       not-applicable  
volumes           ok  
7 entries were displayed.
```

- b. To check the status of the running MetroCluster check operation, use the command:

```
metrocluster operation history show -job-id 38
```

- c. Verify that there are no health alerts:

```
system health alert show
```

2. Prepare the IP switches for the application of the new RCF files.

Resetting the Cisco IP switch to factory defaults

Before installing a new software version and RCFs, you must erase the Cisco switch configuration and perform basic configuration.

You must repeat these steps on each of the IP switches in the MetroCluster IP configuration.

1. Reset the switch to factory defaults:
 - a. Erase the existing configuration: `write erase`
 - b. Reload the switch software: `reload`

The system reboots and enters the configuration wizard. During the boot, if you receive the prompt `Abort Auto Provisioning and continue with normal setup?(yes/no)[n]`, you should respond `yes` to proceed.

c. In the configuration wizard, enter the basic switch settings:

- Admin password
- Switch name
- Out-of-band management configuration
- Default gateway
- SSH service (RSA)

After completing the configuration wizard, the switch reboots.

d. When prompted, enter the user name and password to log in to the switch.

The following example shows the prompts and system responses when configuring the switch. The angle brackets (<<<) show where you enter the information.

```
---- System Admin Account Setup ----
Do you want to enforce secure password standard (yes/no) [y]:y
**<<<**

Enter the password for "admin": password
Confirm the password for "admin": password
---- Basic System Configuration Dialog VDC: 1 ----

This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.

Please register Cisco Nexus3000 Family devices promptly with your
supplier. Failure to register may affect response times for initial
service calls. Nexus3000 devices must be registered to receive
entitled support services.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to
skip the remaining dialogs.
```

You enter basic information in the next set of prompts, including the switch name, management address, and gateway, and select SSH with RSA.

The following configuration will be applied:

```
password strength-check
  switchname IP_switch_A_1
vrf context management
ip route 0.0.0.0/0 10.10.99.1
exit
no feature telnet
ssh key rsa 1024 force
feature ssh
system default switchport
system default switchport shutdown
copp profile strict
interface mgmt0
ip address 10.10.99.10 255.255.255.0
no shutdown
```

Would you like to edit the configuration? (yes/no) [n]:

Use this configuration and save it? (yes/no) [y]:

```
2017 Jun 13 21:24:43 A1 %$ VDC-1 %$ %COPP-2-COPP_POLICY: Control-
Plane is protected with policy copp-system-p-policy-strict.
```

```
[#####] 100%
Copy complete.
```

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
.
.
.
IP_switch_A_1#
```

2. Save the configuration:

```
IP_switch-A-1# copy running-config startup-config
```

3. Reboot the switch and wait for the switch to reload:

```
IP_switch-A-1# reload
```

4. Repeat the previous steps on the other three switches in the MetroCluster IP configuration.

Downloading and installing the Cisco switch NX-OS software

You must download the switch operating system file and RCF file to each switch in the MetroCluster IP configuration.

This task requires file transfer software, such as FTP, TFTP, SFTP, or SCP, to copy the files to the switches.

These steps must be repeated on each of the IP switches in the MetroCluster IP configuration.

You must use the supported switch software version.

NetApp Hardware Universe

1. Download the supported NX-OS software file.

Cisco Software Download

2. Copy the switch software to the switch: `copy sftp://root@server-ip-address/tftpboot/NX-OS-file-name bootflash: vrf management`

In this example, the `nxos.7.0.3.I4.6.bin` file is copied from SFTP server 10.10.99.99 to the local bootflash:

```
IP_switch_A_1# copy sftp://root@10.10.99.99/tftpboot/nxos.7.0.3.I4.6.bin
bootflash: vrf management
root@10.10.99.99's password: password
sftp> progress
Progress meter enabled
sftp> get /tftpboot/nxos.7.0.3.I4.6.bin
/bootflash/nxos.7.0.3.I4.6.bin
Fetching /tftpboot/nxos.7.0.3.I4.6.bin to /bootflash/nxos.7.0.3.I4.6.bin
/tftpboot/nxos.7.0.3.I4.6.bin          100% 666MB 7.2MB/s
01:32
sftp> exit
Copy complete, now saving to disk (please wait)...
```

3. Verify on each switch that the switch NX-OS files are present in each switch's bootflash directory: `dir bootflash:`

The following example shows that the files are present on `IP_switch_A_1`:

```

IP_switch_A_1# dir bootflash:
      .
      .
      .
698629632   Jun 13 21:37:44 2017   nxos.7.0.3.I4.6.bin
      .
      .
      .

Usage for bootflash://sup-local
 1779363840 bytes used
13238841344 bytes free
15018205184 bytes total
IP_switch_A_1#

```

4. Install the switch software: install all nxos bootflash:nxos.version-number.bin

The switch will reload (reboot) automatically after the switch software has been installed.

The following example shows the software installation on IP_switch_A_1:

```

IP_switch_A_1# install all nxos bootflash:nxos.7.0.3.I4.6.bin
Installer will perform compatibility check first. Please wait.
Installer is forced disruptive

Verifying image bootflash:/nxos.7.0.3.I4.6.bin for boot variable "nxos".
[#####] 100% -- SUCCESS

Verifying image type.
[#####] 100% -- SUCCESS

Preparing "nxos" version info using image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS

Preparing "bios" version info using image
bootflash:/nxos.7.0.3.I4.6.bin.
[#####] 100% -- SUCCESS          [#####] 100%
-- SUCCESS

Performing module support checks.          [#####] 100%
-- SUCCESS

Notifying services about system upgrade.   [#####] 100%
-- SUCCESS

```

Compatibility check is done:

Module	bootable	Impact	Install-type	Reason
1	yes	disruptive	reset	default upgrade is not hitless

Images will be upgraded according to following table:

Module	Image	Running-Version (pri:alt)	New-Version	Upg-Required
1	nxos	7.0(3)I4(1)	7.0(3)I4(6)	yes
1	bios	v04.24(04/21/2016)	v04.24(04/21/2016)	no

Switch will be reloaded for disruptive upgrade.

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Performing runtime checks. [#####] 100% --
SUCCESS

Setting boot variables.
[#####] 100% -- SUCCESS

Performing configuration copy.
[#####] 100% -- SUCCESS

Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.
Warning: please do not remove or power off the module at this time.
[#####] 100% -- SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
IP_switch_A_1#

5. Wait for the switch to reload and then log in to the switch.

After the switch has rebooted the login prompt is displayed:

```
User Access Verification
IP_switch_A_1 login: admin
Password:
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
.
MDP database restore in progress.
IP_switch_A_1#

The switch software is now installed.
```

6. Verify that the switch software has been installed: `show version`

The following example shows the output:

```

IP_switch_A_1# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2017, Cisco and/or its affiliates.
All rights reserved.
.
.
.

Software
  BIOS: version 04.24
  NXOS: version 7.0(3)I4(6)   **<<< switch software version**
  BIOS compile time: 04/21/2016
  NXOS image file is: bootflash:///nxos.7.0.3.I4.6.bin
  NXOS compile time: 3/9/2017 22:00:00 [03/10/2017 07:05:18]

Hardware
  cisco Nexus 3132QV Chassis
  Intel(R) Core(TM) i3- CPU @ 2.50GHz with 16401416 kB of memory.
  Processor Board ID FOC20123GPS

  Device name: A1
  bootflash: 14900224 kB
  usb1: 0 kB (expansion flash)

Kernel uptime is 0 day(s), 0 hour(s), 1 minute(s), 49 second(s)

Last reset at 403451 usecs after Mon Jun 10 21:43:52 2017

Reason: Reset due to upgrade
System version: 7.0(3)I4(1)
Service:

plugin
  Core Plugin, Ethernet Plugin
IP_switch_A_1#

```

7. Repeat these steps on the remaining three IP switches in the MetroCluster IP configuration.

Configuring MACsec encryption on Cisco 9336C switches

If desired, you can configure MACsec encryption on the WAN ISL ports that run between the sites. You must configure MACsec after applying the correct RCF file.



MACsec encryption can only be applied to the WAN ISL ports.

Licensing requirements for MACsec

MACsec requires a security license. For a complete explanation of the Cisco NX-OS licensing scheme and how to obtain and apply for licenses, see the [Cisco NX-OS Licensing Guide](#)

Enabling Cisco MACsec Encryption WAN ISLs in MetroCluster IP configurations

You can enable MACsec encryption for Cisco 9336C switches on the WAN ISLs in a MetroCluster IP configuration.

1. Enter the global configuration mode: `configure terminal`

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Enable MACsec and MKA on the device: `feature macsec`

```
IP_switch_A_1(config)# feature macsec
```

3. Copy the running configuration to the startup configuration: `copy running-config startup-config`

```
IP_switch_A_1(config)# copy running-config startup-config
```

Disabling Cisco MACsec Encryption

You might need to disable MACsec encryption for Cisco 9336C switches on the WAN ISLs in a MetroCluster IP configuration.



If you disable encryption, you must also delete your keys.

1. Enter the global configuration mode: `configure terminal`

```
IP_switch_A_1# configure terminal
IP_switch_A_1(config)#
```

2. Disable the MACsec configuration on the device: `macsec shutdown`

```
IP_switch_A_1(config)# macsec shutdown
```



Selecting the no option restores the MACsec feature.

3. Select the interface that you already configured with MACsec.

You can specify the interface type and identity. For an Ethernet port, use `ethernet slot/port`.

```
IP_switch_A_1(config)# interface ethernet 1/15
switch(config-if)#
```

4. Remove the keychain, policy and fallback-keychain configured on the interface to remove the MACsec configuration: `no macsec keychain keychain-name policy policy-name fallback-keychain keychain-name`

```
IP_switch_A_1(config-if)# no macsec keychain kc2 policy abc fallback-
keychain fb_kc2
```

5. Repeat steps 3 and 4 on all interfaces where MACsec is configured.
6. Copy the running configuration to the startup configuration: `copy running-config startup-config`

```
IP_switch_A_1(config)# copy running-config startup-config
```

Configuring a MACsec key chain and keys

For details on configuring a MACsec key chain, see the Cisco documentation for your switch.

Connecting the new NS224 shelf

Steps

1. Install the rail mount kit that came with your shelf by using the installation flyer that came in the kit box.
2. Install and secure the shelf onto the support brackets and rack or cabinet by using the installation flyer.
3. Connect the power cords to the shelf, secure them in with the power cord retainer, and then connect the power cords to different power sources for resiliency.

A shelf powers up when connected to a power source; it does not have power switches. When functioning correctly, a power supply's bicolored LED illuminates green.

4. Set the shelf ID to a number that is unique within the HA pair and across the configuration.
5. Connect the shelf ports in the following order:
 - a. Connect NSM-A, e0a to the switch (Switch-A1 or Switch-B1)
 - b. Connect NSM-B, e0a to the switch (Switch-A2 or Switch-B2)
 - c. Connect NSM-A, e0b to the switch (Switch-A1 or Switch-B1)
 - d. Connect NSM-B, e0b to the switch (Switch-A2 or Switch-B2)
6. Use the cabling layout generated from the **RcfFileGenerator** tool to cable the shelf to the appropriate ports.

Once the new shelf is cabled correctly, ONTAP automatically detects it on the network.

Configure end-to-end encryption in a MetroCluster IP configuration

Beginning with ONTAP 9.15.1, you can configure end-to-end encryption on supported systems to encrypt back-end traffic, such as NVlog and storage replication data, between the sites in a MetroCluster IP configuration.

About this task

- You must be a cluster administrator to perform this task.
- Before you can configure end-to-end encryption, you must [Configure external key management](#).
- Review the supported systems and minimum ONTAP release required to configure end-to-end encryption in a MetroCluster IP configuration:

Minimum ONTAP release	Supported systems
ONTAP 9.17.1	<ul style="list-style-type: none">• AFF A800, AFF C800• AFF A20, AFF A30, AFF C30, AFF A50, AFF C60• AFF A70, AFF A90, AFF A1K, AFF C80• FAS50, FAS70, FAS90
ONTAP 9.15.1	<ul style="list-style-type: none">• AFF A400• AFF C400• FAS8300• FAS8700

Enable end-to-end encryption

Perform the following steps to enable end-to-end encryption.

Steps

1. Verify the health of the MetroCluster configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- b. After the `metrocluster check run` operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

c. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

d. Verify that there are no health alerts:

```
system health alert show
```

2. Verify that external key management is configured on both clusters:

```
security key-manager external show-status
```

3. Enable end-to-end encryption for each DR group:

```
metrocluster modify -is-encryption-enabled true -dr-group-id  
<dr_group_id>
```

Example

```

cluster_A::~*> metrocluster modify -is-encryption-enabled true -dr-group
-id 1
Warning: Enabling encryption for a DR Group will secure NVLog and
Storage
        replication data sent between MetroCluster nodes and have an
impact on
        performance. Do you want to continue? {y|n}: y
[Job 244] Job succeeded: Modify is successful.

```

Repeat this step for each DR group in the configuration.

4. Verify that end-to-end encryption is enabled:

```
metrocluster node show -fields is-encryption-enabled
```

Example

```

cluster_A::~*> metrocluster node show -fields is-encryption-enabled

dr-group-id cluster      node      configuration-state is-encryption-
enabled
-----
1           cluster_A    node_A_1  configured         true
1           cluster_A    node_A_2  configured         true
1           cluster_B    node_B_1  configured         true
1           cluster_B    node_B_2  configured         true
4 entries were displayed.

```

Disable end-to-end encryption

Perform the following steps to disable end-to-end encryption.

Steps

1. Verify the health of the MetroCluster configuration.
 - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::~*> metrocluster check run
```

The operation runs in the background.

b. After the `metrocluster check run` operation completes, run:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show

Component          Result
-----
nodes              ok
lifs               ok
config-replication ok
aggregates        ok
clusters          ok
connections       ok
volumes           ok
7 entries were displayed.
```

c. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

d. Verify that there are no health alerts:

```
system health alert show
```

2. Verify that external key management is configured on both clusters:

```
security key-manager external show-status
```

3. Disable end-to-end encryption on each DR group:

```
metrocluster modify -is-encryption-enabled false -dr-group-id
<dr_group_id>
```

Example

```
cluster_A::*> metrocluster modify -is-encryption-enabled false -dr-group
-id 1
[Job 244] Job succeeded: Modify is successful.
```

Repeat this step for each DR group in the configuration.

4. Verify that end-to-end encryption is disabled:

```
metrocluster node show -fields is-encryption-enabled
```

Example

```
cluster_A::*> metrocluster node show -fields is-encryption-enabled

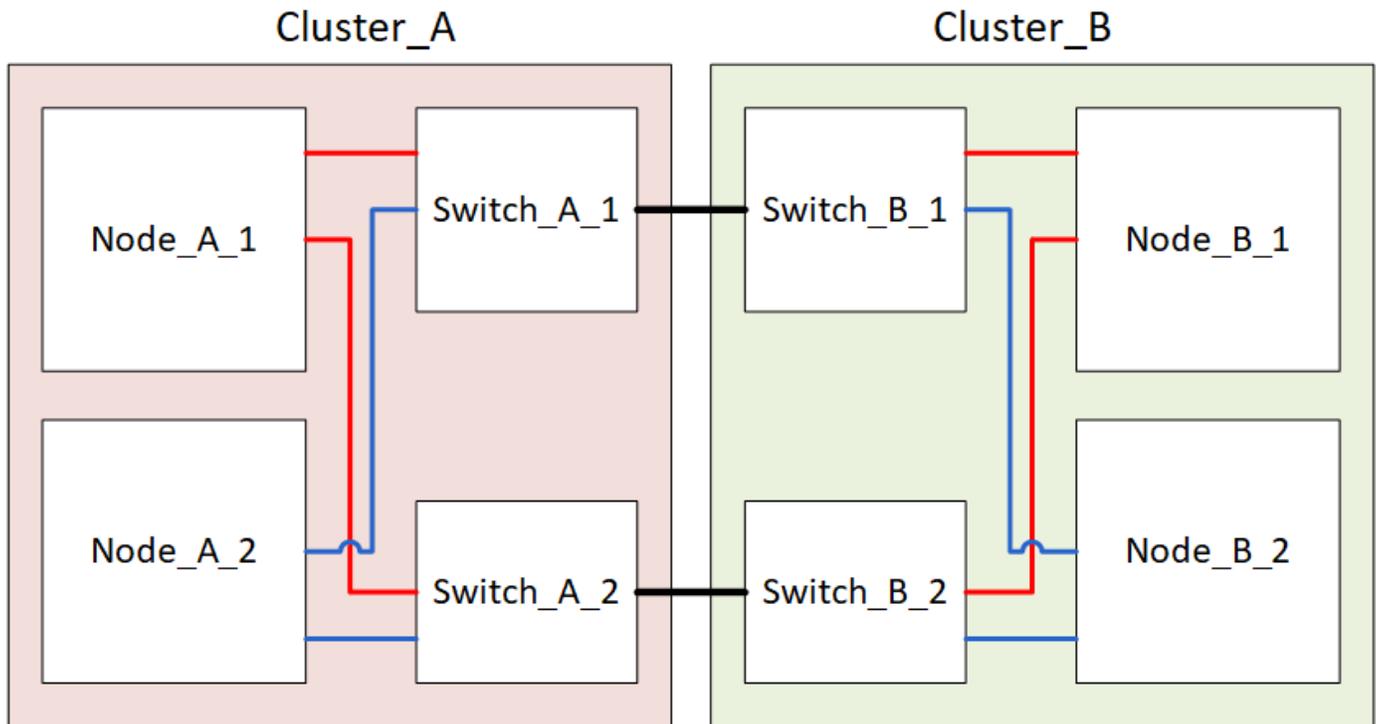
dr-group-id cluster      node      configuration-state is-encryption-
enabled
-----
1           cluster_A    node_A_1  configured         false
1           cluster_A    node_A_2  configured         false
1           cluster_B    node_B_1  configured         false
1           cluster_B    node_B_2  configured         false
4 entries were displayed.
```

Power off and power on a single site in a MetroCluster IP configuration

If you need to perform site maintenance or relocate a single site in a MetroCluster IP configuration, you must know how to power off and power on the site.

If you need to relocate and reconfigure a site (for example, if you need to expand from a four-node to an eight-node cluster), you cannot complete these tasks at the same time. This procedure only covers the steps that are required to perform site maintenance or to relocate a site without changing its configuration.

The following diagram shows a MetroCluster configuration. Cluster_B is powered off for maintenance.



Power off a MetroCluster site

You must power off a site and all of the equipment before site maintenance or relocation can begin.

About this task

All the commands in the following steps are issued from the site that remains powered on.

Steps

1. Before you begin, check that any non-mirrored aggregates at the site are offline.
2. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

3. From the site you want to remain up, implement the switchover:

```
metrocluster switchover
```

```
cluster_A::*> metrocluster switchover
```

The operation can take several minutes to complete.

4. Monitor and verify the completion of the switchover:

```
metrocluster operation show
```

```
cluster_A::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: in-progress
End time: -
Errors:

cluster_A::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: successful
End time: 10/4/2012 19:04:22
Errors: -
```

5. If you have a MetroCluster IP configuration running ONTAP 9.6 or later, wait for the disaster site plexes to come online and the healing operations to automatically complete.

In MetroCluster IP configurations running ONTAP 9.5 or earlier, the disaster site nodes do not automatically boot to ONTAP and the plexes remain offline.

6. Move any volumes and LUNs that belong to unmirrored aggregates offline.

- a. Move the volumes offline.

```
cluster_A::* volume offline <volume name>
```

b. Move the LUNs offline.

```
cluster_A::* lun offline lun_path <lun_path>
```

7. Move unmirrored aggregates offline: `storage aggregate offline`

```
cluster_A*::> storage aggregate offline -aggregate <aggregate-name>
```

8. Depending on your configuration and ONTAP version, identify and move offline affected plexes that are located at the disaster site (Cluster_B).

You should move the following plexes offline:

- Non-mirrored plexes residing on disks located at the disaster site.

If you do not move the non-mirrored plexes at the disaster site offline, an outage might occur when the disaster site is later powered off.

- Mirrored plexes residing on disks located at the disaster site for aggregate mirroring. After they are moved offline, the plexes are inaccessible.

a. Identify the affected plexes.

Plexes that are owned by nodes at the surviving site consist of Pool1 disks. Plexes that are owned by nodes at the disaster site consist of Pool0 disks.

```

Cluster_A::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate      plex  status          is-online pool
-----
Node_B_1_aggr0 plex0 normal,active true      0
Node_B_1_aggr0 plex1 normal,active true      1

Node_B_2_aggr0 plex0 normal,active true      0
Node_B_2_aggr0 plex5 normal,active true      1

Node_B_1_aggr1 plex0 normal,active true      0
Node_B_1_aggr1 plex3 normal,active true      1

Node_B_2_aggr1 plex0 normal,active true      0
Node_B_2_aggr1 plex1 normal,active true      1

Node_A_1_aggr0 plex0 normal,active true      0
Node_A_1_aggr0 plex4 normal,active true      1

Node_A_1_aggr1 plex0 normal,active true      0
Node_A_1_aggr1 plex1 normal,active true      1

Node_A_2_aggr0 plex0 normal,active true      0
Node_A_2_aggr0 plex4 normal,active true      1

Node_A_2_aggr1 plex0 normal,active true      0
Node_A_2_aggr1 plex1 normal,active true      1
14 entries were displayed.

Cluster_A::>

```

The affected plexes are those that are remote to cluster A. The following table shows whether the disks are local or remote relative to cluster A:

Node	Disks in pool	Should the disks be set offline?	Example of plexes to be moved offline
Node_A_1 and Node_A_2	Disks in pool 0	No. Disks are local to cluster A.	-
	Disks in pool 1	Yes. Disks are remote to cluster A.	Node_A_1_aggr0/plex4 Node_A_1_aggr1/plex1 Node_A_2_aggr0/plex4 Node_A_2_aggr1/plex1

Node_B_1 and Node_B_2	Disks in pool 0	Yes. Disks are remote to cluster A.	Node_B_1_aggr1/plex0 Node_B_1_aggr0/plex0 Node_B_2_aggr0/plex0 Node_B_2_aggr1/plex0
	Disks in pool 1	No. Disks are local to cluster A.	-

b. Move the affected plexes offline:

```
storage aggregate plex offline
```

```
storage aggregate plex offline -aggregate Node_B_1_aggr0 -plex plex0
```



Perform this step for all plexes that have disks that are remote to Cluster_A.

9. Persistently offline the ISL switch ports according to the switch type.

10. Halt the nodes by running the following command on each node:

```
node halt -inhibit-takeover true -skip-lif-migration true -node <node-name>
```

11. Power off the equipment at the disaster site.

You must power off the following equipment in the order shown:

- Storage controllers - the storage controllers should currently be at the `LOADER` prompt, you must power them off completely.
- MetroCluster IP switches
- Storage shelves

Relocating the powered-off site of the MetroCluster

After the site is powered off, you can begin maintenance work. The procedure is the same whether the MetroCluster components are relocated within the same data center or relocated to a different data center.

- The hardware should be cabled in the same way as the previous site.
- If the Inter-Switch Link (ISL) speed, length, or number has changed, they all need to be reconfigured.

Steps

1. Verify that the cabling for all components is carefully recorded so that it can be correctly reconnected at the new location.
2. Physically relocate all the hardware, storage controllers, IP switches, and storage shelves.
3. Configure the ISL ports and verify the intersite connectivity.
 - a. Power on the IP switches.



Do **not** power up any other equipment.

4. Use tools on the switches (as they are available) to verify the intersite connectivity.



You should only proceed if the links are correctly configured and stable.

5. Disable the links again if they are found to be stable.

Powering on the MetroCluster configuration and returning to normal operation

After maintenance has been completed or the site has been moved, you must power on the site and reestablish the MetroCluster configuration.

About this task

All the commands in the following steps are issued from the site that you power on.

Steps

1. Power on the switches.

You should power on the switches first. They might have been powered on during the previous step if the site was relocated.

- a. Reconfigure the Inter-Switch Link (ISL) if required or if this was not completed as part of the relocation.
 - b. Enable the ISL if fencing was completed.
 - c. Verify the ISL.
2. Power on the storage controllers and wait until you see the `LOADER` prompt. The controllers must not be fully booted.

If auto boot is enabled, press `Ctrl+C` to stop the controllers from automatically booting.



Don't power up the shelves before you power up the controllers. This prevents the controllers from an unintended boot into ONTAP.

3. Power on the shelves, allowing enough time for them to power on completely.
4. Verify that the storage is visible from maintenance mode.

- a. Boot into maintenance mode:

```
boot_ontap maint
```

- b. Verify that the storage is visible from the surviving site.
- c. Verify that the local and remote storage is visible from the node in maintenance mode:

```
disk show -v
```

5. Halt the nodes:

```
halt
```

6. Reestablish the MetroCluster configuration.

Follow the instructions in [Verifying that your system is ready for a switchback](#) to perform healing and switchback operations according to your MetroCluster configuration.

Powering off an entire MetroCluster IP configuration

You must power off the entire MetroCluster IP configuration and all of the equipment before maintenance or relocation can begin.



Beginning with ONTAP 9.8, the **storage switch** command is replaced with **system switch**. The following steps show the **storage switch** command, but if you are running ONTAP 9.8 or later, the **system switch** command is preferred.

1. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Confirm that the MetroCluster configuration and operational mode are normal.

```
metrocluster show
```

b. Run the following command:

```
metrocluster interconnect show
```

c. Confirm connectivity to the disks by entering the following command on any one of the MetroCluster nodes:

```
run local sysconfig -v
```

d. Run the following command:

```
storage port show
```

e. Run the following command:

```
storage switch show
```

f. Run the following command:

```
network interface show
```

g. Run the following command:

```
network port show
```

h. Run the following command:

```
network device-discovery show
```

i. Perform a MetroCluster check:

```
metrocluster check run
```

j. Display the results of the MetroCluster check:

```
metrocluster check show
```

k. Run the following command:

metrocluster configuration-settings interface show

2. If necessary, disable AUSO by modifying the AUSO Failure Domain to

auso-disabled

```
cluster_A_site_A::*>metrocluster modify -auto-switchover-failure-domain
auso-disabled
```



In a MetroCluster IP configuration, the AUSO Failure Domain is already set to 'auso-disabled' unless the configuration is configured with ONTAP Mediator.

3. Verify the change using the command

metrocluster operation show

```
cluster_A_site_A::*> metrocluster operation show
Operation: modify
State: successful
Start Time: 4/25/2020 20:20:36
End Time: 4/25/2020 20:20:36
Errors: -
```

4. Halt the nodes:

halt

```
system node halt -node node1_SiteA -inhibit-takeover true -ignore-quorum
-warnings true
```

5. Power off the following equipment at the site:
 - Storage controllers
 - MetroCluster IP switches
 - Storage shelves
6. Wait for thirty minutes and then power on all storage shelves, MetroCluster IP switches, and storage controllers.
7. After the controllers are powered on, verify the MetroCluster configuration from both sites.

To verify the configuration, repeat step 1.

8. Perform power cycle checks.
 - a. Verify that all sync-source SVMs are online:

vserver show

- b. Start any sync-source SVMs that are not online:

```
vserver start
```

Maintenance procedures for all MetroCluster configurations

Replacing a shelf nondisruptively in a stretch MetroCluster configuration

You can replace disk shelves without disruption in a stretch MetroCluster configuration with a fully populated disk shelf or a disk shelf chassis and transfer components from the shelf you are removing.

The disk shelf model you are installing must meet the storage system requirements specified in the [Hardware Universe](#), which includes supported shelf models, supported disk drive types, the maximum number of disk shelves in a stack, and supported ONTAP versions.

Steps

1. Properly ground yourself.
2. Identify all aggregates and volumes that have disks from the loop that contains the shelf you are replacing and make note of the affected plex name.

Either node might contain disks from the loop of the affected shelf and host aggregates or host volumes.

3. Choose one of the following two options based on the replacement scenario you are planning.
 - If you are replacing a complete disk shelf, including the shelf chassis, disks, and I/O modules (IOM), take the corresponding action as described in the table below:

Scenario	Action
The affected plex contains fewer disks from the affected shelf.	Replace the disks one-by-one on the affected shelf with spares from another shelf.  You can take the plex offline after completing the disk replacement.
The affected plex contains more disks than are in the affected shelf.	Move the plex offline and then delete the plex.
The affected plex has any disk from the affected shelf.	Move the plex offline but do not delete it.

- If you are replacing only the disk shelf chassis and no other components, perform the following steps:
 - a. Offline the affected plexes from the controller where they are hosted:

```
aggregate offline
```

- b. Verify that the plexes are offline:

```
aggregate status -r
```

4. Identify the controller SAS ports to which the affected shelf loop is connected and disable the SAS ports on both site controllers:

```
storage port disable -node node_name -port SAS_port
```

The affected shelf loop is connected to both sites.

5. Wait for ONTAP to recognize that the disk is missing.

- a. Verify that the disk is missing:

```
sysconfig -a or sysconfig -r
```

6. Turn off the power switch on the disk shelf.
7. Unplug all power cords from the disk shelf.
8. Make a record of the ports from which you unplug the cables so that you can cable the new disk shelf in the same way.
9. Unplug and remove the cables connecting the disk shelf to the other disk shelves or the storage system.
10. Remove the disk shelf from the rack.

To make the disk shelf lighter and easier to maneuver, remove the power supplies and IOM. If you will be installing a disk shelf chassis, also remove the disk drives or carriers. Otherwise, avoid removing disk drives or carriers if possible because excessive handling can cause internal drive damage.

11. Install and secure the replacement disk shelf onto the support brackets and rack.
12. If you installed a disk shelf chassis, reinstall power supplies and IOM.
13. Reconfigure the stack of disk shelves by connecting all cables to the replacement disk shelf ports exactly as they were configured on the disk shelf that you removed.
14. Turn on the power to the replacement disk shelf and wait for the disk drives to spin up.
15. Change the disk shelf ID to a unique ID from 0 through 98.
16. Enable any SAS ports that you previously disabled .
 - a. Wait for ONTAP to recognize that the disks are inserted.
 - b. Verify that the disks are inserted:

```
sysconfig -a or sysconfig -r
```

17. If you are replacing the complete disk shelf (disk shelf chassis, disks, IOM), perform the following steps:



If you are replacing only the disk shelf chassis and no other components, go to Step 19.

- a. Determine whether disk auto assignment is enabled (on).

```
storage disk option modify -autoassign
```

Disk assignment will occur automatically.

- b. If disk auto assignment is not enabled, assign disk ownership manually.

18. Move the plexes back online:

```
aggregate online plex name
```

19. Recreate any plexes that were deleted by mirroring the aggregate.

20. Monitor the plexes as they begin resynchronizing:

```
aggregate status -r <aggregate name>
```

21. Verify that the storage system is functioning as expected:

```
system health alert show
```

When to migrate root volumes to a new destination

You might need to move root volumes to another root aggregate within a two-node or four-node MetroCluster configuration.

Migrating root volumes within a two-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a two-node MetroCluster configuration, you should refer to [How to move mroot to a new root aggregate in a 2-node Clustered MetroCluster with Switchover](#). This procedure shows you how to non-disruptively migrate the root volumes during a MetroCluster switchover operation. This procedure is slightly different than the procedure used on a four-node configuration.

Migrating root volumes within a four-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a four-node MetroCluster configuration, you can use the [system node migrate-root](#) command while meeting the following requirements.

- You can use `system node migrate-root` to move root aggregates within a four-node MetroCluster configuration.
- All root aggregates must be mirrored.
- You can add new shelves on both sites with smaller drives to host the root aggregate.
- You must check the drive limits that the platform supports before attaching new drives.

[NetApp Hardware Universe](#)

- If you move the root aggregate to smaller drives, you need to accommodate the minimum root volume size of the platform to ensure all core files are saved.



The four-node procedure can also be applied to an eight-node configuration.

Moving a metadata volume in MetroCluster configurations

You can move a metadata volume from one aggregate to another aggregate in a MetroCluster configuration. You might want to move a metadata volume when the source aggregate is decommissioned or unmirrored, or for other reasons that make the aggregate ineligible.

- You must have cluster administrator privileges to perform this task.

- The target aggregate must be mirrored and should not be in the degraded state.
- The available space in the target aggregate must be larger than the metadata volume that you are moving.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Identify the metadata volume that should be moved:

```
volume show MDV_CRS*
```

```
Cluster_A::*> volume show MDV_CRS*
Vserver   Volume                Aggregate             State                Type                Size
Available Used%
-----
Cluster_A
          MDV_CRS_14c00d4ac9f311e7922800a0984395f1_A
                Node_A_1_aggr1
                        online                RW                10GB
9.50GB    5%
Cluster_A
          MDV_CRS_14c00d4ac9f311e7922800a0984395f1_B
                Node_A_2_aggr1
                        online                RW                10GB
9.50GB    5%
Cluster_A
          MDV_CRS_15035e66c9f311e7902700a098439625_A
                Node_B_1_aggr1
                        -                RW                -
-         -
Cluster_A
          MDV_CRS_15035e66c9f311e7902700a098439625_B
                Node_B_2_aggr1
                        -                RW                -
-         -
4 entries were displayed.

Cluster_A::>
```

3. Identify an eligible target aggregate:

```
metrocluster check config-replication show-aggregate-eligibility
```

The following command identifies the aggregates in cluster_A that are eligible to host metadata volumes:

```
Cluster_A::*> metrocluster check config-replication show-aggregate-
eligibility
```

```
Aggregate Hosted Config Replication Vols Host Addl Vols Comments
-----
-----
Node_A_1_aggr0 - false Root Aggregate
Node_A_2_aggr0 - false Root Aggregate
Node_A_1_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A true -
Node_A_2_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_B true -
Node_A_1_aggr2 - true
Node_A_2_aggr2 - true
Node_A_1_Aggr3 - false Unable to determine available space of aggregate
Node_A_1_aggr5 - false Unable to determine mirror configuration
Node_A_2_aggr6 - false Mirror configuration does not match requirement
Node_B_1_aggr4 - false NonLocal Aggregate
```



In the previous example, Node_A_1_aggr2 and Node_A_2_aggr2 are eligible.

4. Start the volume move operation:

```
volume move start -vserver svm_name -volume metadata_volume_name -destination
-aggregate destination_aggregate_name
```

The following command moves metadata volume MDV_CRS_14c00d4ac9f311e7922800a0984395f1 from aggregate Node_A_1_aggr1 to aggregate Node_A_1_aggr2:

```
Cluster_A::*> volume move start -vserver svm_cluster_A -volume
MDV_CRS_14c00d4ac9f311e7922800a0984395f1
-destination-aggregate aggr_cluster_A_02_01

Warning: You are about to modify the system volume
"MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A". This may cause
severe
performance or stability problems. Do not proceed unless
directed to
do so by support. Do you want to proceed? {y|n}: y
[Job 109] Job is queued: Move
"MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" in Vserver
"svm_cluster_A" to aggregate "aggr_cluster_A_02_01".
Use the "volume move show -vserver svm_cluster_A -volume
MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" command to view the status
of this operation.
```

5. Verify the state of the volume move operation:

```
volume move show -volume vol_constituent_name
```

6. Return to the admin privilege level:

```
set -privilege admin
```

Renaming a cluster in MetroCluster configurations

Renaming a cluster in a MetroCluster configuration involves making the changes, and then verifying on both the local and remote clusters that the change took effect correctly.

Steps

1. View the cluster names using the

```
metrocluster node show
```

command:

```
cluster_1::*> metrocluster node show
DR
Group Cluster Node          Configuration  DR
-----
State          Mirroring Mode
-----
1      cluster_1
      node_A_1      configured   enabled   normal
      node_A_2      configured   enabled   normal
      cluster_2
      node_B_1      configured   enabled   normal
      node_B_2      configured   enabled   normal
4 entries were displayed.
```

2. Rename the cluster:

```
cluster identity modify -name new_name
```

In the following example, the `cluster_1` cluster is renamed `cluster_A`:

```
cluster_1::*> cluster identity modify -name cluster_A
```

3. Verify on the local cluster that the renamed cluster is running normally:

```
metrocluster node show
```

In the following example, the newly renamed `cluster_A` is running normally:

```

cluster_A::*> metrocluster node show
DR
Group Cluster Node          Configuration  DR
State          Mirroring Mode
-----
-----
1      cluster_A
      node_A_1      configured   enabled   normal
      node_A_2      configured   enabled   normal
      cluster_2
      node_B_1      configured   enabled   normal
      node_B_2      configured   enabled   normal
4 entries were displayed.

```

4. Rename the remote cluster:

```
cluster peer modify-local-name -name cluster_2 -new-name cluster_B
```

In the following example, cluster_2 is renamed cluster_B:

```

cluster_A:::> cluster peer modify-local-name -name cluster_2 -new-name
cluster_B

```

5. Verify on the remote cluster that the local cluster was renamed and is running normally:

```
metrocluster node show
```

In the following example, the newly renamed cluster_B is running normally:

```

cluster_B::*> metrocluster node show
DR
Group Cluster Node          Configuration  DR
State          Mirroring Mode
-----
-----
1      cluster_B
      node_B_1      configured   enabled   normal
      node_B_2      configured   enabled   normal
      cluster_A
      node_A_1      configured   enabled   normal
      node_A_2      configured   enabled   normal
4 entries were displayed.

```

6. Repeat these steps for each cluster that you want to rename.

Verify the health of a MetroCluster configuration

Learn how to verify that the MetroCluster components are healthy.

About this task

- In MetroCluster IP and FC configurations, you can use the CLI to run health check commands and verify the state of the MetroCluster components.
- In MetroCluster IP configurations running ONTAP 9.8 or later, you can also use ONTAP System Manager to monitor and troubleshoot health check alerts.

Steps

Verify the health of the MetroCluster configuration depending on whether you are using the CLI or System Manager.

CLI

Use the follow steps to check the health of a MetroCluster configuration using the CLI.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

2. After the `metrocluster check run` operation completes, display the results:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok

7 entries were displayed.

3. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

4. Verify that there are no health alerts:

```
system health alert show
```

ONTAP System Manager (MetroCluster IP only)

Beginning with ONTAP 9.8, System Manager monitors the health of MetroCluster IP configurations and helps you identify and correct problems that might occur.

System Manager periodically checks the health of your MetroCluster IP configuration. When you view the MetroCluster section in the Dashboard, usually the message is "MetroCluster systems are healthy."

However, when a problem occurs, the message will show the number of events. You can click on this message and view the results of the health check for the following components:

- Node
- Network Interface
- Tier (Storage)
- Cluster
- Connection
- Volume
- Configuration Replication

The **Status** column identifies which components have problems, and the **Details** column suggests how to correct the problem.

Steps

1. In System Manager, select **Dashboard**.
2. View the message in the **MetroCluster** section:
 - a. If the message indicates that your MetroCluster configuration is healthy, and the connections between the clusters and the ONTAP Mediator are healthy (shown with check marks), then you have no problems to correct.
 - b. If the message lists the number of events, or the connections have gone down (shown with an "X"), then continue to the next step.
3. Click the message that shows the number of events.

The MetroCluster Health Report displays.

4. Troubleshoot the problems that appear in the report using the suggestions in the **Details** column.
5. When all the problems have been corrected, click **Check MetroCluster Health**.



You should perform all your troubleshooting tasks before running the check because the MetroCluster Health Check uses an intensive amount of resources.

The MetroCluster Health Check runs in the background. You can work on other tasks while you wait for it to finish.

Where to find additional information

You can learn more about configuring, operating, and monitoring a MetroCluster configuration in NetApp's extensive documentation.

Information	Subject
-------------	---------

MetroCluster documentation	<ul style="list-style-type: none"> • All MetroCluster information
NetApp MetroCluster Solution Architecture and Design	<ul style="list-style-type: none"> • A technical overview of the MetroCluster configuration and operation. • Best practices for MetroCluster configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the MetroCluster IP configuration • Configuring the MetroCluster in ONTAP
NetApp Documentation: Product Guides and Resources	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance
MetroCluster Tiebreaker Software installation and configuration	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems

Transition from MetroCluster FC to MetroCluster IP

Choose your transition procedure

When transitioning to a MetroCluster IP configuration, you must have a combination of supported platform models.

You should also ensure that the MetroCluster IP platform is an appropriate size for the load that you are transitioning from the MetroCluster FC configuration to the MetroCluster IP configuration.

Supported platform combinations

- The transition procedures all require ONTAP 9.8 or later unless stated otherwise in the notes or as required by an individual platform.
- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, if you have an eight-node configuration, all eight nodes must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.



- Do not exceed any object limits of the 'lower' of the platforms in the combination. Apply the lower object limit of the two platforms.
- If the target platform limits are lower than the MetroCluster limits, you must reconfigure the MetroCluster to be at, or below, the target platform limits before you add the new nodes.
- Refer to the [Hardware universe](#) for platform limits.

Supported AFF and FAS transition combinations

The following tables show the supported platform combinations. You can transition from platforms in the first column to platforms listed as supported in the columns to the right, as indicated by the colored table cells.

For example, transitioning from a MetroCluster FC configuration consisting of AFF8060 controller modules to an IP configuration consisting of AFF A400 controller modules is supported.

The tables are split into two groups:

- **Group 1** shows combinations for transitions to AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.
- **Group 2** shows combinations for transitions to AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

The following notes apply to both groups:

- Note 1: This platform combination requires ONTAP 9.11.1 or later.
- Note 2: You must have a 40GbE interface for the local cluster interfaces on the FC nodes.
- Note 3: You must have a 100GbE interface for the local cluster interfaces on the FC nodes.

AFF and FAS combinations group 1

Review the supported combinations for transitions to AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.

AFF and FAS		Target MetroCluster IP platform									
		AFF A150	AFF A20	FAS2750 AFF A220	FAS500f AFF C250 AFF A250	FAS50	AFF C30 AFF A30	FAS8200 AFF A300	AFF A320	FAS8300 AFF C400 AFF A400	FAS8700
Source MetroCluster FC platform	FAS8020 AFF8020 FAS8040 AFF8040										
	FAS8060 AFF8060 FAS8080 AFF8080										
	FAS8200 AFF A300				Note 1						
	FAS8300 AFF A400										
	FAS9000 AFF A700										
	FAS9500 AFF A900										

AFF and FAS combinations group 2

Review the supported combinations for transitions to AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

AFF and FAS		Target MetroCluster IP platform									
		AFF C60	AFF A50	FAS70	FAS9000 AFF A700	AFF A70	AFF C800 AFF A800	FAS9500 AFF A900	AFF C80	FAS90 AFF A90	AFF A1K
Source MetroCluster FC platform	FAS8020 AFF8020 FAS8040 AFF8040										
	FAS8060 AFF8060 FAS8080 AFF8080										
	FAS8200 AFF A300										
	FAS8300 AFF A400										
	FAS9000 AFF A700			Note 2	Note 2	Note 2	Note 2	Note 2	Note 2	Note 2	Note 2
	FAS9500 AFF A900							Note 3	Note 3	Note 3	Note 3

Supported ASA transition platform combinations

The following table shows the supported platform combinations for ASA systems.

Source MetroCluster FC platform	Target MetroCluster IP platform	Supported?
ASA A400	ASA A400	Yes
	ASA A900	No
ASA A900	ASA A400	No
	ASA A900	Yes

Choose your transition procedure

You must select a transition procedure depending on your existing MetroCluster FC configuration.

A transition procedure replaces the back-end FC switch fabric or FC-VI connection with an IP switch network. The exact procedure depends on your starting configuration.

The original platforms and FC switches (if present) are retired at the end of the transition procedure.

Starting configuration	Disruptive or nondisruptive	Requirements	Procedure
Four or eight node	Nondisruptive	New storage shelves are required on new platforms. After the transition is completed, old controllers, shelves and disks are removed from the cluster.	Link to procedure Note: This procedure supports the following FC to IP transitions: <ul style="list-style-type: none"> • From a four-node MetroCluster FC configuration to a four-node MetroCluster IP configuration • From an eight-node MetroCluster FC configuration to an eight-node MetroCluster IP configuration
Two node	Disruptive	New storage shelves are supported on both original and new platforms.	Link to procedure
Two node	Disruptive	New storage shelves are supported on both original and new platforms. Old storage shelves must be retired.	Link to procedure
Two node	Disruptive	Old storage shelves are not supported on new platforms. Old storage shelves must be retired.	Link to procedure

Transition nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later)

Transitioning nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later)

You can perform nondisruptive transitions of workloads and data from an existing MetroCluster FC configuration to a new MetroCluster IP configuration.

Beginning with ONTAP 9.13.1, this procedure is supported in MetroCluster IP configurations in which the MetroCluster and the drive shelves are connected to the same IP switches (a shared storage switch configuration).

Beginning with ONTAP 9.13.1, you can perform a nondisruptive transition of workloads and data from an existing eight-node MetroCluster FC configuration to a new eight-node MetroCluster IP configuration. You use this procedure to transition one four-node FC DR group, remove the empty FC DR group, and then repeat the procedure for the second FC DR group.

Beginning with ONTAP 9.8, you can perform a nondisruptive transition of workloads and data from an existing four-node MetroCluster FC configuration to a new four-node MetroCluster IP configuration. After you complete the transition, you can expand to an eight-node MetroCluster IP configuration if required. See [Expand a MetroCluster IP configuration](#).

- This procedure is nondisruptive.

The MetroCluster configuration can continue to serve data during the operation.

- This procedure applies only to four-node and eight-node MetroCluster FC configurations.

If you have a two-node MetroCluster FC configuration, see [Choosing your transition procedure](#).

- This procedure describes the steps required to transition one four-node FC DR group. If you have an eight-node configuration (two FC DR groups), you must repeat the entire procedure for each FC DR group.



When you add or remove a DR group as part of this procedure, you must verify that the DR group removal or addition was successful before adding or removing another DR group.

- You must meet all requirements and follow all steps in the procedure.

Important information if you are adding an older platform model

The following guidance is for an uncommon scenario where you need to add an older platform model (platforms released before ONTAP 9.15.1) to an existing MetroCluster configuration that contains a newer platform model (platforms released in ONTAP 9.15.1 or later). For an eight-node FC to IP transition, this guidance applies if you have transitioned your first FC DR group to a platform model introduced in ONTAP 9.15.1 or later and plan to transition the second DR group to a platform introduced before ONTAP 9.15.1.

If your existing MetroCluster configuration contains a platform that uses **shared cluster/HA ports** (platforms released in ONTAP 9.15.1 or later), you cannot add a platform that uses **shared MetroCluster/HA ports** (platforms released before ONTAP 9.15.1) without upgrading all nodes in the configuration to ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later.



Adding an older platform model that uses **shared/MetroCluster HA ports** to a MetroCluster containing a newer platform model that uses **shared cluster/HA ports** is an uncommon scenario and most combinations are not affected.

Use the following table to verify whether your combination is affected. If your existing platform is listed in the first column, and the platform you want to add to the configuration is listed in the second column, all nodes in the configuration must be running ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later to add the new DR group.

If your existing MetroCluster contains..		And the platform you're adding is...		Then...
An AFF system using shared cluster/HA ports : <ul style="list-style-type: none"> • AFF A20 • AFF A30 • AFF C30 • AFF A50 • AFF C60 • AFF C80 • AFF A70 • AFF A90 • AFF A1K 	A FAS system using shared cluster/HA ports : <ul style="list-style-type: none"> • FAS50 • FAS70 • FAS90 	An AFF system using shared MetroCluster/HA ports : <ul style="list-style-type: none"> • AFF A150, ASA A150 • AFF A220 • AFF C250, ASA C250 • AFF A250, ASA A250 • AFF A300 • AFF A320 • AFF C400, ASA C400 • AFF A400, ASA A400 • AFF A700 • AFF C800, ASA C800 • AFF A800, ASA A800 • AFF A900, ASA A900 	A FAS system using shared MetroCluster/HA ports : <ul style="list-style-type: none"> • FAS2750 • FAS500f • FAS8200 • FAS8300 • FAS8700 • FAS9000 • FAS9500 	Before you add the new platform to your existing MetroCluster configuration, upgrade all nodes in the existing and new configuration to ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later.

Prepare for transition from a MetroCluster FC to a MetroCluster IP configuration

Enable console logging

Enable console logging on your devices before performing this task.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the

duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Requirements for nondisruptive FC-to-IP transition

Before starting the transition process, verify that the configuration meets the requirements.

You can perform a nodisruptive FC-to-IP transition if your configuration meets the following requirements:

- If you have an eight-node configuration, all nodes are running ONTAP 9.13.1 or later.
- If you have a four-node configuration, all nodes are running ONTAP 9.8 or later.
- The existing and new platforms are a supported combination for transition.

[Supported platforms for nondisruptive transition](#)

- Your configuration supports a switched cluster configuration.

[Hardware Universe](#)



If you are using shared storage MetroCluster switches, you can only transition to a four-node MetroCluster IP configuration. Transitioning to an eight-node MetroCluster IP configuration using shared storage MetroCluster switches is not supported. After you complete the transition to a four-node MetroCluster IP configuration, you can [expand to an eight-node MetroCluster IP configuration](#).

- Your configuration meets all requirements and is cabled as described in the following *MetroCluster Installation and Configuration* procedures.

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

How transition impacts the MetroCluster hardware components

After completing the transition procedure, key components of the existing MetroCluster configuration have been replaced or reconfigured.

• **Controller modules**

The existing controller modules are replaced by new controller modules. The existing controller modules are decommissioned at the end of the transition procedures.

• **Storage shelves**

Data is moved from the old shelves to the new shelves. The old shelves are decommissioned at the end of

the transition procedures.

- **MetroCluster (back-end) and cluster switches**

The back-end switch functionality is replaced by the IP switch fabric. If the MetroCluster FC configuration included FC switches and FC-to-SAS bridges, they are decommissioned at the end of this procedure.

If the MetroCluster FC configuration used cluster switches for the cluster interconnect, in some cases they can be reused to provide the back-end IP switch fabric. Reused cluster switches must be reconfigured with platform and switch-specific RCFs. procedures.

If the MetroCluster FC configuration did not use cluster switches, new IP switches are added to provide the backend switch fabric.

[Considerations for IP switches](#)

- **Cluster peering network**

The existing customer-provided cluster peering network can be used for the new MetroCluster IP configuration. Cluster peering is configured on the MetroCluster IP nodes as part of the transition procedure.

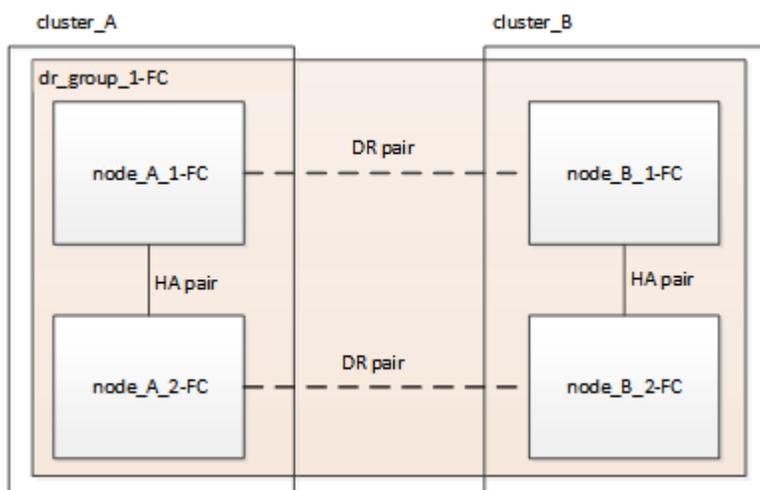
Workflow for nondisruptive MetroCluster transition

You must follow the specific workflow to ensure a successful nondisruptive transition. Choose the workflow for your configuration:

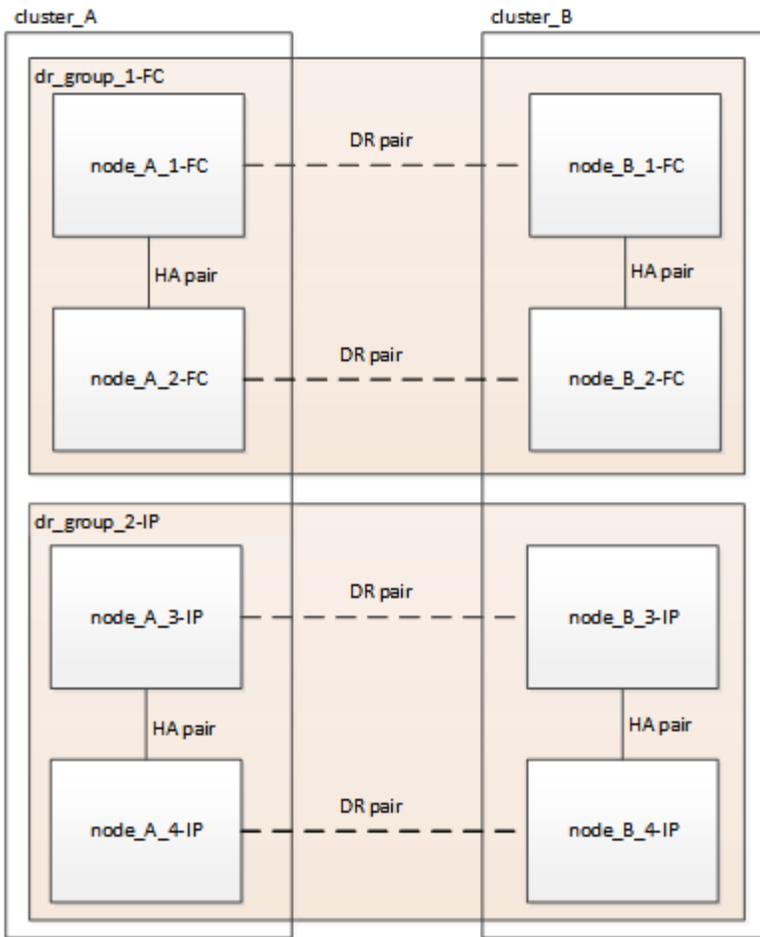
- [Four-node FC configuration transition workflow](#)
- [Eight-node FC configuration transition workflow](#)

Four-node FC configuration transition workflow

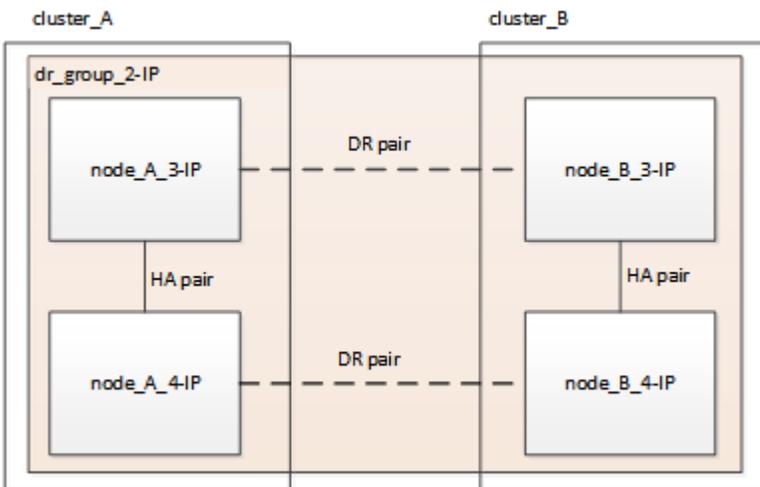
The transition process begins with a healthy four-node MetroCluster FC configuration.



The new MetroCluster IP nodes are added as a second DR group.

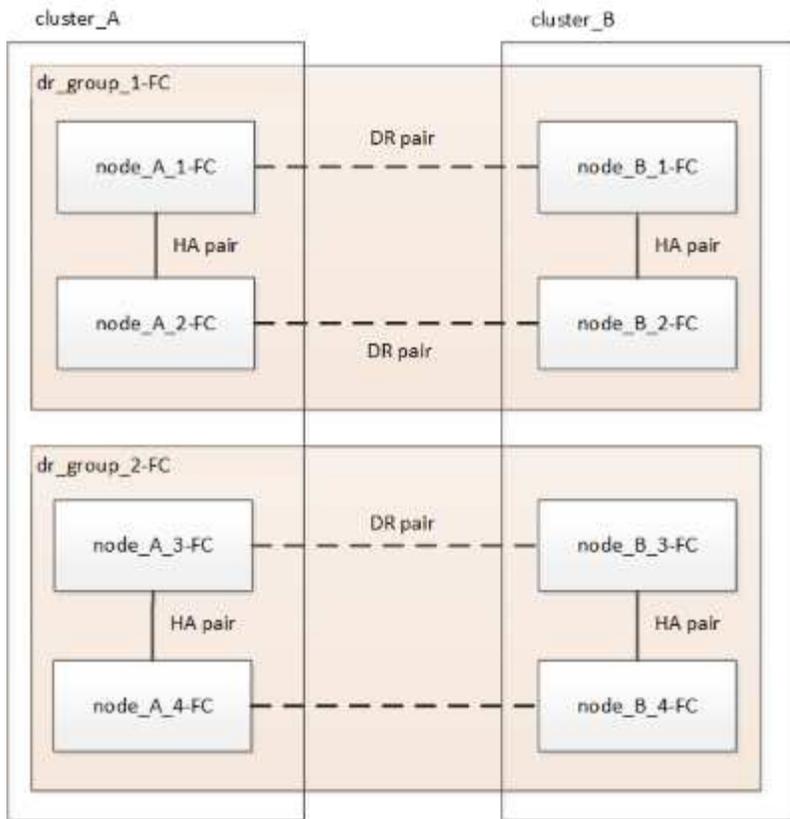


Data is transferred from the old DR group to the new DR group, and then the old nodes and their storage are removed from the configuration and decommissioned. The process ends with a four-node MetroCluster IP configuration.

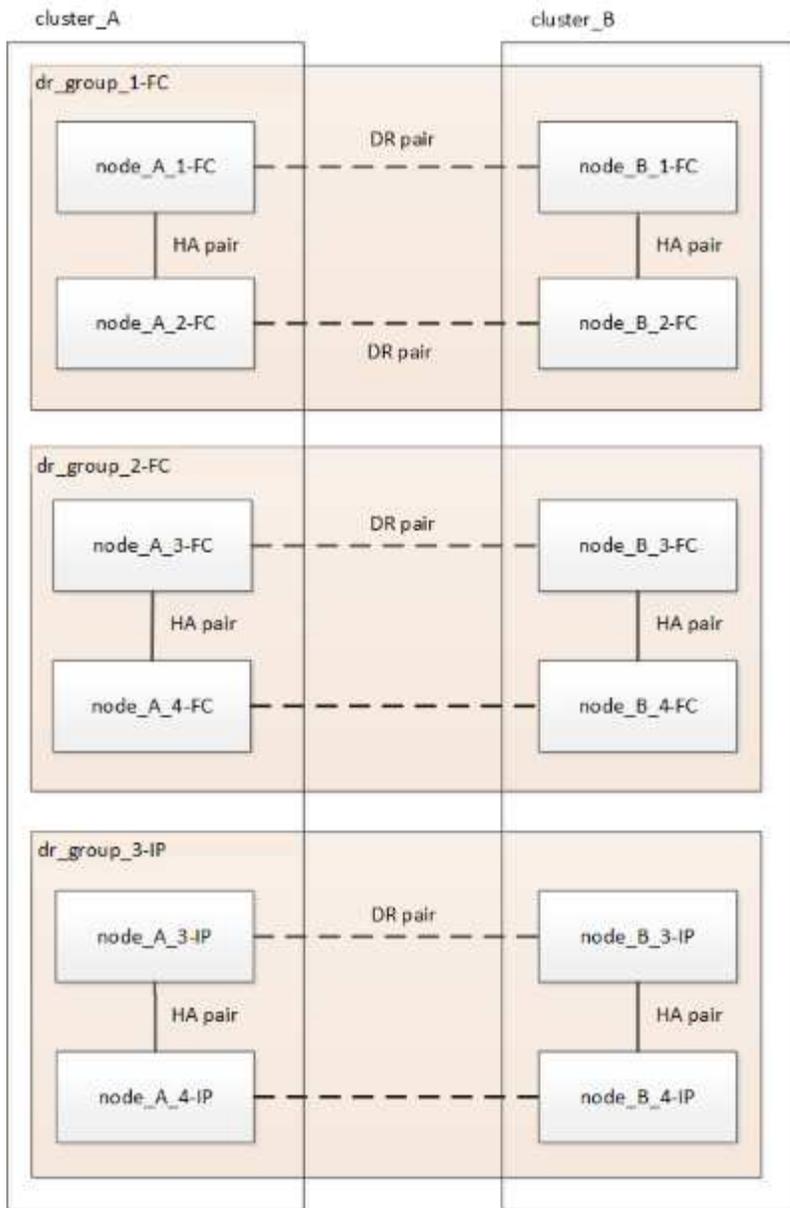


Eight-node FC configuration transition workflow

The transition process begins with a healthy eight-node MetroCluster FC configuration.



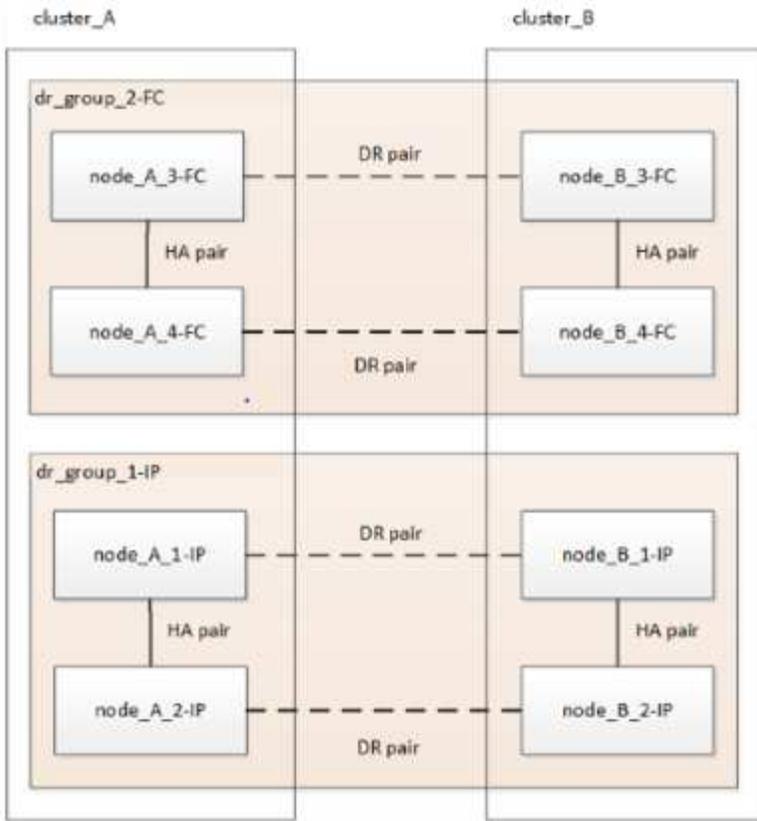
The new MetroCluster IP nodes are added as a third DR group.



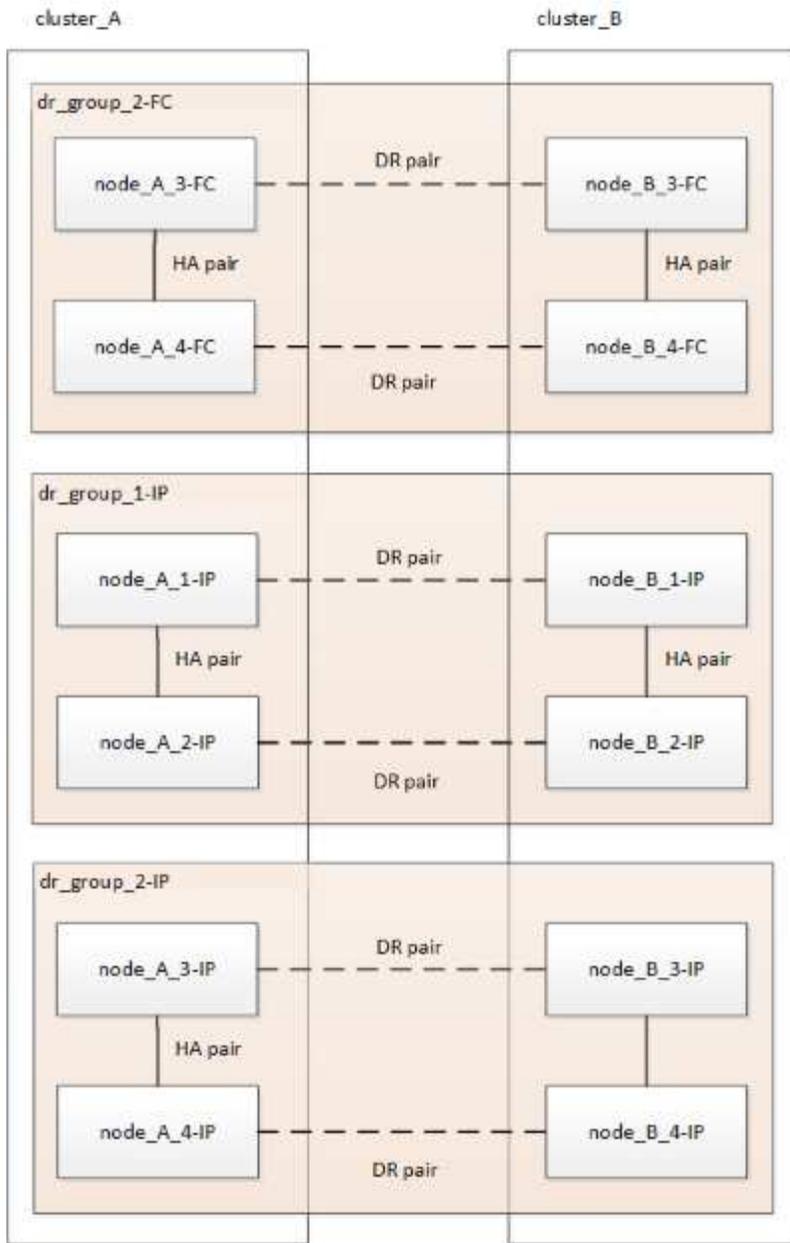
Data is transferred from DR_group_1-FC to DR_group_1-IP, and then the old nodes and their storage are removed from the configuration and decommissioned.



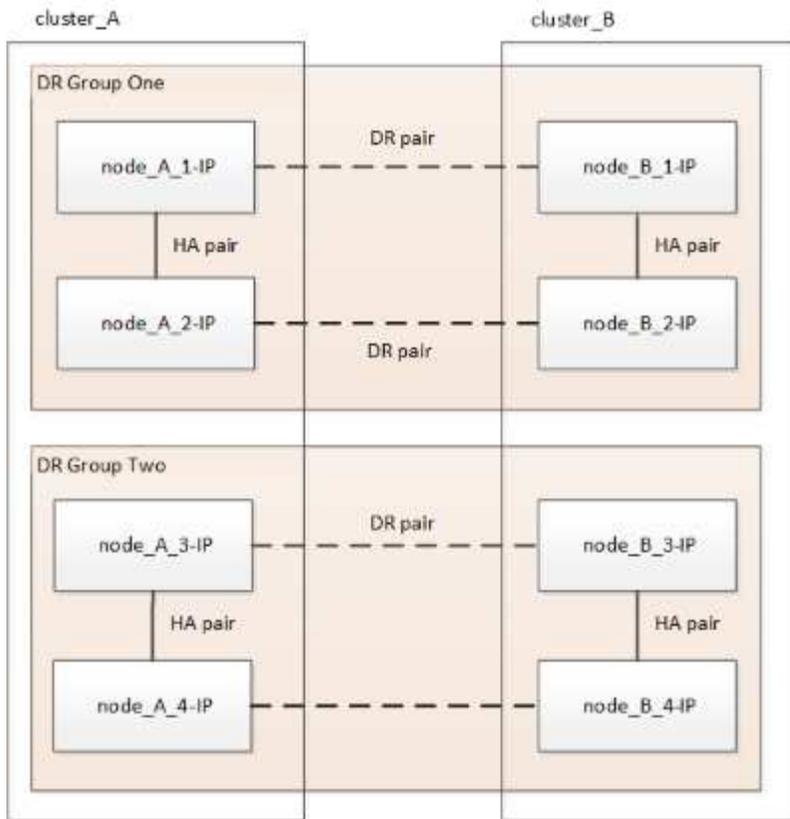
If you want to transition from an eight-node FC configuration to a four-node IP configuration, you must transition all the data in DR_group_1-FC and DR_group_2-FC to the new IP DR group (DR_group_1-IP). You can then decommission both FC DR groups. After the FC DR groups have been removed, the process ends with a four-node MetroCluster IP configuration.



Add the remaining MetroCluster IP nodes to the existing MetroCluster configuration. Repeat the process to transfer data from the DR_group_2-FC nodes to the DR_group_2-IP nodes.

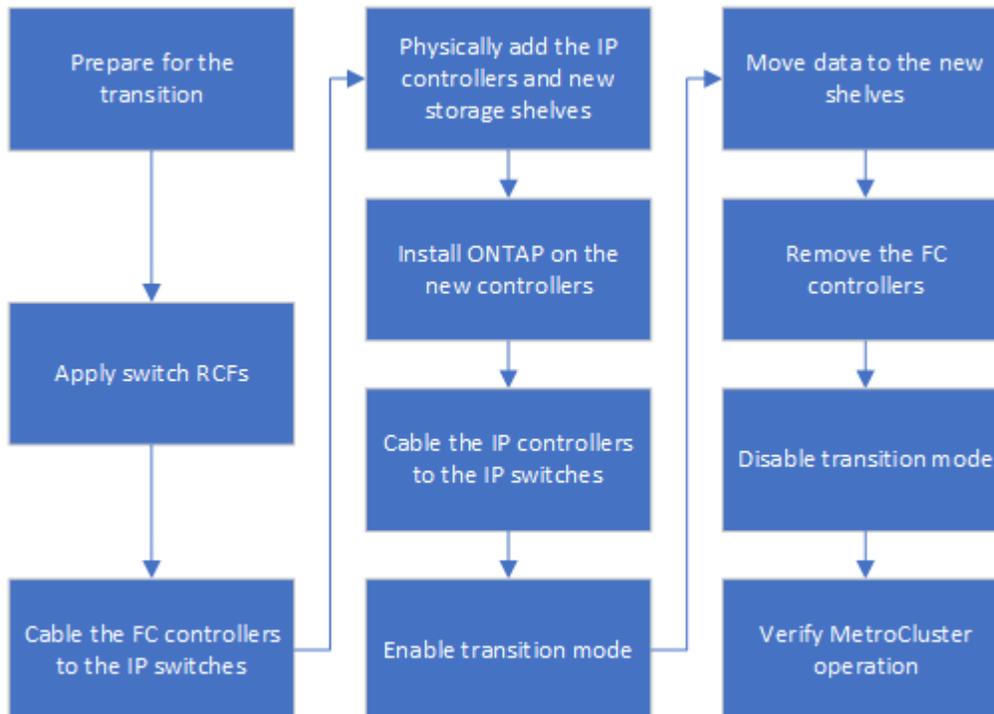


After removing DR_group_2-FC, the process ends with an eight-node MetroCluster IP configuration.



Transition process workflow

You will use the following workflow to transition the MetroCluster configuration.



Considerations for IP switches

You must ensure the IP switches are supported. If the existing switch model is supported

by both the original MetroCluster FC configuration and the new MetroCluster IP configuration, you can reuse the existing switches.

Supported switches

You must use NetApp-provided switches.

- The use of MetroCluster-compliant switches (switches that are not validated and provided by NetApp) is not supported for transition.
- The IP switches must be supported as a cluster switch by both the MetroCluster FC configuration and the MetroCluster IP configuration.
- The IP switches can be reused in the new MetroCluster IP configuration if the MetroCluster FC is a switched cluster and the IP cluster switches are supported by the MetroCluster IP configuration.
- New IP switches are usually used in the following cases:
 - The MetroCluster FC is a switchless cluster, so new switches are required.
 - The MetroCluster FC is a switched cluster but the existing IP switches are not supported in the MetroCluster IP configuration.
 - You want to use different switches for the MetroCluster IP configuration.



If you are using shared storage MetroCluster switches, you can only transition to a four-node MetroCluster IP configuration. Transitioning to an eight-node MetroCluster IP configuration using shared storage MetroCluster switches is not supported. After you complete the transition to a four-node MetroCluster IP configuration, you can [expand to an eight-node MetroCluster IP configuration](#).

See the [Hardware Universe](#) for information on platform models and switch support.

Switchover, healing, and switchback operations during nondisruptive transition

Depending on the stage of the transition process, the MetroCluster switchover, healing, and switchback operations use either the MetroCluster FC or MetroCluster IP workflow.

The following table shows what workflows are used at different stages of the transition process. In some stages, switchover and switchback are not supported.

- In the MetroCluster FC workflow, the switchover, healing, and switchback steps are those used by a MetroCluster FC configuration.
- In the MetroCluster IP workflow, the switchover, healing, and switchback steps are those used by a MetroCluster IP configuration.
- In the unified workflow, when both the FC and IP nodes are configured, the steps depend on whether NSO or USO is performed. The details are shown in the table.

For information on the MetroCluster FC and IP workflows for switchover, healing, and switchback, see [Understanding MetroCluster data protection and disaster recovery](#).



Automatic unplanned switchover is not available during the transition process.

Stage of transition	Negotiated switchover uses this workflow...	Unplanned switchover uses this workflow...
Before the MetroCluster IP nodes have joined the cluster	MetroCluster FC	MetroCluster FC
After the MetroCluster IP nodes have joined the cluster, before the <code>metrocluster configure</code> command is performed	Not supported	MetroCluster FC
After the <code>metrocluster configure</code> command has been issued. Volume move can be in progress.	Unified: All remote site nodes remain up and healing is done automatically	Unified: <ul style="list-style-type: none"> • Mirrored aggregates owned by the MetroCluster FC node are mirrored if storage is accessible, all others are degraded after switchover. • All remote site nodes are able to boot up. • The <code>heal aggregate</code> and <code>heal root</code> commands must be run manually.
The MetroCluster FC nodes have been unconfigured.	Not supported	MetroCluster IP
The <code>cluster unjoin</code> command has been performed on the MetroCluster FC nodes.	MetroCluster IP	MetroCluster IP

Alert messages and tool support during transition

You may notice alert messages during transition. These alerts can be safely ignored. Also, some tools are not available during transition.

- ARS may alert during transition.

These alerts can be ignored and should disappear once the transition has finished.

- OnCommand Unified Manager may alert during transition.

These alerts can be ignored and should disappear once the transition has finished.

- Config Advisor is not supported during transition.
- System Manager is not supported during transition.

Example naming in this procedure

This procedure uses example names throughout to identify the DR groups, nodes, and switches involved.

DR groups	cluster_A at site_A	cluster_B at site_B
dr_group_1-FC	<ul style="list-style-type: none"> • node_A_1-FC • node_A_2-FC 	<ul style="list-style-type: none"> • node_B_1-FC • node_B_2-FC
dr_group_2-IP	<ul style="list-style-type: none"> • node_A_3-IP • node_A_4-IP 	<ul style="list-style-type: none"> • node_B_3-IP • node_B_4-IP
Switches	<p>Initial switches (if fabric-attached configuration:)</p> <ul style="list-style-type: none"> • switch_A_1-FC • switch_A_2-FC <p>MetroCluster IP switches:</p> <ul style="list-style-type: none"> • switch_A_1-IP • switch_A_2-IP 	<p>Initial switches (if fabric-attached configuration:)</p> <ul style="list-style-type: none"> • switch_B_1-FC • switch_B_2-FC <p>MetroCluster IP switches:</p> <ul style="list-style-type: none"> • switch_B_1-IP • switch_B_2-IP

Transition from MetroCluster FC to MetroCluster IP configurations

Verifying the health of the MetroCluster configuration

You must verify the health and connectivity of the MetroCluster configuration prior to performing the transition

1. Verify the operation of the MetroCluster configuration in ONTAP:
 - a. Check whether the system is multipathed: `node run -node node-name sysconfig -a`
 - b. Check for any health alerts on both clusters: `system health alert show`
 - c. Confirm the MetroCluster configuration and that the operational mode is normal: `metrocluster show`
 - d. Perform a MetroCluster check: `metrocluster check run`
 - e. Display the results of the MetroCluster check: `metrocluster check show`
 - f. Check for any health alerts on the switches (if present): `storage switch show`
 - g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
2. Verify that the cluster is healthy: `cluster show`

```

cluster_A::> cluster show
Node           Health  Eligibility  Epsilon
-----
node_A_1_FC   true   true         false
node_A_2_FC   true   true         false

cluster_A::>

```

3. Verify that all cluster ports are up: `network port show -ipspace cluster`

```

cluster_A::> network port show -ipspace cluster

Node: node_A_1_FC

Port           IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
Admin/Oper    Status
-----
e0a            Cluster      Cluster      up   9000    auto/10000 healthy
e0b            Cluster      Cluster      up   9000    auto/10000 healthy

Node: node_A_2_FC

Port           IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
Admin/Oper    Status
-----
e0a            Cluster      Cluster      up   9000    auto/10000 healthy
e0b            Cluster      Cluster      up   9000    auto/10000 healthy

4 entries were displayed.

cluster_A::>

```

4. Verify that all cluster LIFs are up and operational: `network interface show -vserver cluster`

Each cluster LIF should display "true" for "Is Home" and "up/up" for "Status Admin/Oper".

```
cluster_A::> network interface show -vserver cluster
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
	-----	-----	-----	-----	-----
Cluster					
true	node_A-1_FC_clus1	up/up	169.254.209.69/16	node_A-1_FC	e0a
true	node_A_1_FC_clus2	up/up	169.254.49.125/16	node_A_1_FC	e0b
true	node_A_2_FC_clus1	up/up	169.254.47.194/16	node_A_2_FC	e0a
true	node_A_2_FC_clus2	up/up	169.254.19.183/16	node_A_2_FC	e0b

```
4 entries were displayed.
```

```
cluster_A::>
```

5. Verify that auto-revert is enabled on all cluster LIFs: `network interface show -vserver Cluster -fields auto-revert`

```

cluster_A::> network interface show -vserver Cluster -fields auto-revert

          Logical
Vserver   Interface      Auto-revert
-----
Cluster
          node_A_1_FC_clus1
                        true
          node_A_1_FC_clus2
                        true
          node_A_2_FC_clus1
                        true
          node_A_2_FC_clus2
                        true

          4 entries were displayed.

cluster_A::>

```

Removing the existing configuration from the Tiebreaker or other monitoring software

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the Tiebreaker or other software prior to transition.

1. Remove the existing MetroCluster configuration from the Tiebreaker software.

Removing MetroCluster configurations

2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

Generating and applying RCFs to the new IP switches

If you are using new IP switches for the MetroCluster IP configuration, you must configure the switches with a custom RCF file.

This task is required if you are using new switches.

If you are using existing switches, proceed to [Moving the local cluster connections](#).

1. Install and rack the new IP switches.
2. Prepare the IP switches for the application of the new RCF files.

Follow the steps in the section for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
 - [Reset the Cisco IP switch to factory defaults](#)
 - [Reset the NVIDIA IP SN2100 switch to factory defaults](#)
3. Update the firmware on the switch to a supported version, if necessary.
 4. Use the RCF generator tool to create the RCF file depending on your switch vendor and the platform models, and then update the switches with the file.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom IP RCF files](#)
- [Download and install the Cisco IP RCF files](#)
- [Download and install the NVIDIA RCF files](#)

Move the local cluster connections

Move the MetroCluster FC configuration's cluster interfaces to the IP switches.

Step 1: Move the cluster connections on the MetroCluster FC nodes

Move the cluster connections on the MetroCluster FC nodes to the IP switches. The steps you follow depend on whether you're using existing IP switches or new IP switches.

About this task

- You perform this task on both MetroCluster sites.

Which connections to move

The following task assumes a controller module is using two ports for the cluster connections. Some controller module models use four or more ports for the cluster connection. In this example, the ports are divided into two groups, alternating ports between the two groups.

The following table shows the example ports used in this task.

Number of cluster connections on the controller module	Group A ports	Group B ports
Two	e0a	e0b
Four	e0a, e0c	e0b, e0d

- Group A ports connect to local switch switch_x_1-IP.
- Group B ports connect to local switch switch_x_2-IP.

The following table shows which switch ports the FC nodes connect to. For the Broadcom BES-53248 switch, the port usage depends on the model of the MetroCluster IP nodes.

Switch model	MetroCluster IP node model	Switch port(s)	Connects to
--------------	----------------------------	----------------	-------------

Cisco 3132Q-V	Any	5, 6	Local cluster interface on FC node
Cisco 9336C-FX2 (12-port)	Any	3,4, or 11,12 Note: To use switch ports 11 and 12, you must select two speed modes.	Local cluster interface on FC node
Cisco 3232C, or 9336C-FX2 (36-port)	Any	5, 6, or 13, 14 Note: To use switch ports 13 and 14, you must select two speed modes.	Local cluster interface on FC node
Cisco 9336C-FX2 shared (36-port)	Any	3,4, or 11,12 Note: To use switch ports 11 and 12, you must select two speed modes.	Local cluster interface on FC node
Broadcom BES-53248	FAS500f/A250	1 - 6	Local cluster interface on FC node
	FAS8200/A300	3, 4, 9, 10, 11, 12	Local cluster interface on FC node
	FAS8300/A400/FAS8700	1 - 6	Local cluster interface on FC node
NVIDIA SN2100	Any	5,6, or 11,12 Note: To use switch ports 11 and 12, you must select two speed modes.	Local cluster interface on FC node

Move the local cluster connections when using new IP switches

If you are using new IP switches, you physically move the existing MetroCluster FC nodes' cluster connections to the new switches.

Steps

1. Move the MetroCluster FC node group A cluster connections to the new IP switches.

Use the ports described in [Which connections to move](#).

- a. Disconnect all the group A ports from the switch, or, if the MetroCluster FC configuration was a switchless cluster, disconnect them from the partner node.
- b. Disconnect the group A ports from node_A_1-FC and node_A_2-FC.

c. Connect the group A ports of node_A_1-FC to the switch ports for the FC node on switch_A_1-IP

d. Connect the group A ports of node_A_2-FC to the switch ports for the FC node on switch_A_1-IP

2. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```

```
cluster_A::~*> network port show -ipspace Cluster

Node: node_A_1-FC

Port          IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
-----
Admin/Oper    Status
-----
e0a           Cluster      Cluster          up  9000    auto/10000 healthy
e0b           Cluster      Cluster          up  9000    auto/10000 healthy

Node: node_A_2-FC

Port          IPspace      Broadcast Domain Link MTU      Speed(Mbps) Health
-----
Admin/Oper    Status
-----
e0a           Cluster      Cluster          up  9000    auto/10000 healthy
e0b           Cluster      Cluster          up  9000    auto/10000 healthy

4 entries were displayed.

cluster_A::~*>
```

3. Verify that your inter-site Inter-Switch Links (ISLs) are up and the port-channels are operational:

```
show interface brief
```

In the following example, ISL ports “Eth1/15” to “Eth1/20” are configured as “Po10” for the remote site link and “Eth1/7” to “Eth1/8” are configured as “Po1” for the local cluster ISL. The state of “Eth1/15” to “Eth1/20”, “Eth1/7” to “Eth1/8”, “Po10”, and “Po1” should be 'up'.

```
IP_switch_A_1# show interface brief

-----
Port    VRF      Status  IP Address      Speed  MTU
-----
mgmt0  --      up      100.10.200.20  1000  1500
-----

Ethernet  VLAN  Type Mode  Status  Reason  Speed
Port
```

```

Interface                                                    Ch #
-----
...

Eth1/7              1      eth  trunk  up      none      100G(D)
1
Eth1/8              1      eth  trunk  up      none      100G(D)
1
...

Eth1/15             1      eth  trunk  up      none      100G(D)
10
Eth1/16             1      eth  trunk  up      none      100G(D)
10
Eth1/17             1      eth  trunk  up      none      100G(D)
10
Eth1/18             1      eth  trunk  up      none      100G(D)
10
Eth1/19             1      eth  trunk  up      none      100G(D)
10
Eth1/20             1      eth  trunk  up      none      100G(D)
10

-----
-----
Port-channel VLAN  Type Mode  Status  Reason          Speed  Protocol
Interface
-----
-----
Po1              1      eth  trunk  up      none            a-100G(D) lacp
Po10             1      eth  trunk  up      none            a-100G(D) lacp
Po11             1      eth  trunk  down    No operational auto(D)  lacp
members

IP_switch_A_1#

```

4. Verify that all interfaces display true in the “Is Home” column:

```
network interface show -vserver cluster
```

This might take several minutes to complete.

```

cluster_A::*> network interface show -vserver cluster

          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
Cluster
          node_A_1_FC_clus1
          up/up      169.254.209.69/16  node_A_1_FC  e0a
true
          node_A_1-FC_clus2
          up/up      169.254.49.125/16  node_A_1-FC  e0b
true
          node_A_2-FC_clus1
          up/up      169.254.47.194/16  node_A_2-FC  e0a
true
          node_A_2-FC_clus2
          up/up      169.254.19.183/16  node_A_2-FC  e0b
true

4 entries were displayed.

cluster_A::*>

```

5. Perform the above steps on both nodes (node_A_1-FC and node_A_2-FC) to move the group B ports of the cluster interfaces.
6. Repeat the above steps on the partner cluster "cluster_B".

Move the local cluster connections when reusing existing IP switches

If you are reusing existing IP switches, you update firmware, reconfigure the switches with the correct reference configuration files (RCFs) and move the connections to the correct ports one switch at a time.

About this task

This task is required only if the FC nodes are connected to existing IP switches and you are reusing the switches.

Steps

1. Disconnect the local cluster connections that connect to switch_A_1_IP
 - a. Disconnect the group A ports from the existing IP switch.
 - b. Disconnect the ISL ports on switch_A_1_IP.

You can see the Installation and Setup instructions for the platform to see the cluster port usage.

[AFF A320 systems: Installation and setup](#)

[AFF A220/FAS2700 Systems Installation and Setup Instructions](#)

[AFF A800 Systems Installation and Setup Instructions](#)

[AFF A300 Systems Installation and Setup Instructions](#)

[FAS8200 Systems Installation and Setup Instructions](#)

2. Reconfigure switch_A_1_IP using RCF files generated for your platform combination and transition.

Follow the steps in the procedure for your switch vendor from *MetroCluster IP Installation and Configuration*:

[MetroCluster IP installation and configuration](#)

- a. If required, download and install the new switch firmware.

You should use the latest firmware that the MetroCluster IP nodes support.

- [Download and install the Broadcom switch EFOS software](#)
- [Download and install the Cisco switch NX-OS software](#)
- [Download and install the NVIDIA Cumulus software](#)

- b. Prepare the IP switches for the application of the new RCF files.

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

- c. Download and install the IP RCF file depending on your switch vendor.

- [Download and installing the Broadcom IP RCF files](#)
- [Download and installing the Cisco IP RCF files](#)
- [Download and install the NVIDIA RCF files](#)

3. Reconnect the group A ports to switch_A_1_IP.

Use the ports described in [Which connections to move](#).

4. Verify that all cluster ports are up:

```
network port show -ip space cluster
```

```
Cluster-A::*> network port show -ipspace cluster
```

```
Node: node_A_1_FC
```

Port	IPspace	Broadcast	Domain	Link	MTU	Speed(Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster		up	9000	auto/10000	healthy
e0b	Cluster	Cluster		up	9000	auto/10000	healthy

```
Node: node_A_2_FC
```

Port	IPspace	Broadcast	Domain	Link	MTU	Speed(Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster		up	9000	auto/10000	healthy
e0b	Cluster	Cluster		up	9000	auto/10000	healthy

```
4 entries were displayed.
```

```
Cluster-A::*>
```

5. Verify that all interfaces are on their home port:

```
network interface show -vserver Cluster
```

```

Cluster-A::*> network interface show -vserver Cluster

          Logical      Status      Network      Current
Current Is
Vserver   Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
-----
Cluster
          node_A_1_FC_clus1
          up/up      169.254.209.69/16  node_A_1_FC  e0a
true
          node_A_1_FC_clus2
          up/up      169.254.49.125/16  node_A_1_FC  e0b
true
          node_A_2_FC_clus1
          up/up      169.254.47.194/16  node_A_2_FC  e0a
true
          node_A_2_FC_clus2
          up/up      169.254.19.183/16  node_A_2_FC  e0b
true

4 entries were displayed.

Cluster-A::*>

```

6. Repeat all the previous steps on switch_A_2_IP.
7. Reconnect the local cluster ISL ports.
8. Repeat the above steps at site_B for switch B_1_IP and switch B_2_IP.
9. Connect the remote ISLs between the sites.

Step 2: Verify that the cluster connections are moved and the cluster is healthy

To ensure that there is proper connectivity and that the configuration is ready to proceed with the transition process, verify that the cluster connections are moved correctly, the cluster switches are recognized and the cluster is healthy.

Steps

1. Verify that all cluster ports are up and running:

```
network port show -ipSPACE Cluster
```

```
Cluster-A::*> network port show -ipspace Cluster
```

```
Node: Node-A-1-FC
```

Port	IPspace	Broadcast	Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster		up	9000	auto/10000	healthy
e0b	Cluster	Cluster		up	9000	auto/10000	healthy

```
Node: Node-A-2-FC
```

Port	IPspace	Broadcast	Domain	Link	MTU	Speed (Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster		up	9000	auto/10000	healthy
e0b	Cluster	Cluster		up	9000	auto/10000	healthy

```
4 entries were displayed.
```

```
Cluster-A::*>
```

2. Verify that all interfaces are on their home port:

```
network interface show -vserver Cluster
```

This might take several minutes to complete.

The following example shows that all interfaces show true in the “Is Home” column.

```
Cluster-A::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
	-----	-----	-----	-----	-----
Cluster					
true	Node-A-1_FC_clus1	up/up	169.254.209.69/16	Node-A-1_FC	e0a
true	Node-A-1-FC_clus2	up/up	169.254.49.125/16	Node-A-1-FC	e0b
true	Node-A-2-FC_clus1	up/up	169.254.47.194/16	Node-A-2-FC	e0a
true	Node-A-2-FC_clus2	up/up	169.254.19.183/16	Node-A-2-FC	e0b

```
4 entries were displayed.
```

```
Cluster-A::*>
```

3. Verify that both the local IP switches are discovered by the nodes:

```
network device-discovery show -protocol cdp
```

```
Cluster-A::*> network device-discovery show -protocol cdp
```

Node/ Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform

Node-A-1-FC				
	/cdp			
	e0a	Switch-A-3-IP	1/5/1	N3K-
C3232C				
	e0b	Switch-A-4-IP	0/5/1	N3K-
C3232C				
Node-A-2-FC				
	/cdp			
	e0a	Switch-A-3-IP	1/6/1	N3K-
C3232C				
	e0b	Switch-A-4-IP	0/6/1	N3K-
C3232C				

```
4 entries were displayed.
```

```
Cluster-A::*>
```

4. On the IP switch, verify that the MetroCluster IP nodes have been discovered by both local IP switches:

```
show cdp neighbors
```

You must perform this step on each switch.

This example shows how to verify the nodes are discovered on Switch-A-3-IP.

```
(Switch-A-3-IP)# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge  
S - Switch, H - Host, I - IGMP, r - Repeater,  
V - VoIP-Phone, D - Remotely-Managed-Device,  
s - Supports-STP-Dispute
```

Device-ID ID	Local Intrfce	Hldtme	Capability	Platform	Port
Node-A-1-FC	Eth1/5/1	133	H	FAS8200	e0a
Node-A-2-FC	Eth1/6/1	133	H	FAS8200	e0a
Switch-A-4-IP (FDO220329A4)	Eth1/7	175	R S I s	N3K-C3232C	Eth1/7
Switch-A-4-IP (FDO220329A4)	Eth1/8	175	R S I s	N3K-C3232C	Eth1/8
Switch-B-3-IP (FDO220329B3)	Eth1/20	173	R S I s	N3K-C3232C	
Eth1/20					
Switch-B-3-IP (FDO220329B3)	Eth1/21	173	R S I s	N3K-C3232C	
Eth1/21					

```
Total entries displayed: 4
```

```
(Switch-A-3-IP)#
```

This example shows how to verify that the nodes are discovered on Switch-A-4-IP.

```
(Switch-A-4-IP)# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge  
S - Switch, H - Host, I - IGMP, r - Repeater,  
V - VoIP-Phone, D - Remotely-Managed-Device,  
s - Supports-STP-Dispute
```

Device-ID ID	Local Intrfce	Hldtme	Capability	Platform	Port
Node-A-1-FC	Eth1/5/1	133	H	FAS8200	e0b
Node-A-2-FC	Eth1/6/1	133	H	FAS8200	e0b
Switch-A-3-IP (FDO220329A3)	Eth1/7	175	R S I s	N3K-C3232C	Eth1/7
Switch-A-3-IP (FDO220329A3)	Eth1/8	175	R S I s	N3K-C3232C	Eth1/8
Switch-B-4-IP (FDO220329B4)	Eth1/20	169	R S I s	N3K-C3232C	
Eth1/20					
Switch-B-4-IP (FDO220329B4)	Eth1/21	169	R S I s	N3K-C3232C	
Eth1/21					

```
Total entries displayed: 4
```

```
(Switch-A-4-IP)#
```

Preparing the MetroCluster IP controllers

You must prepare the four new MetroCluster IP nodes and install the correct ONTAP version.

This task must be performed on each of the new nodes:

- node_A_1-IP
- node_A_2-IP
- node_B_1-IP
- node_B_2-IP

In these steps, you clear the configuration on the nodes and clear the mailbox region on new drives.

1. Rack the new controllers for the MetroCluster IP configuration.

The MetroCluster FC nodes (node_A_x-FC and node_B_x-FC) remain cabled at this time.

2. Cable the MetroCluster IP nodes to the IP switches as shown in the [Cabling the IP switches](#).
3. Configure the MetroCluster IP nodes using the following sections:

- a. [Gather required information](#)
 - b. [Restore system defaults on a controller module](#)
 - c. [Verify the ha-config state of components](#)
 - d. [Manually assign drives for pool 0 \(ONTAP 9.4 and later\)](#)
4. From Maintenance mode, issue the halt command to exit Maintenance mode, and then issue the boot_ontap command to boot the system and get to cluster setup.

Do not complete the cluster wizard or node wizard at this time.

5. Repeat these steps on the other MetroCluster IP nodes.

Configure the MetroCluster for transition

To prepare the configuration for transition you add the new nodes to the existing MetroCluster configuration and then move data to the new nodes.

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

“maintenance-window-in-hours” specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

2. Repeat the command on the partner cluster.

Enabling transition mode and disabling cluster HA

You must enable the MetroCluster transition mode to allow the old and new nodes to operate together in the MetroCluster configuration, and disable cluster HA.

1. Enable transition:
 - a. Change to the advanced privilege level:

```
set -privilege advanced
```

- b. Enable transition mode:

```
metrocluster transition enable -transition-mode non-disruptive
```



Run this command on one cluster only.

```
cluster_A::~* > metrocluster transition enable -transition-mode non-
disruptive

Warning: This command enables the start of a "non-disruptive"
MetroCluster
          FC-to-IP transition. It allows the addition of hardware for
another DR
          group that uses IP fabrics, and the removal of a DR group
that uses FC
          fabrics. Clients will continue to access their data during a
non-disruptive transition.

          Automatic unplanned switchover will also be disabled by this
command.
Do you want to continue? {y|n}: y

cluster_A::~* >
```

c. Return to the admin privilege level:

```
set -privilege admin
```

2. Verify that transition is enabled on both the clusters.

```
cluster_A::> metrocluster transition show-mode
Transition Mode

non-disruptive

cluster_A::~* >

cluster_B::~* > metrocluster transition show-mode
Transition Mode

non-disruptive

Cluster_B::>
```

3. Disable cluster HA.



You must run this command on both clusters.

```
cluster_A::*> cluster ha modify -configured false
```

```
Warning: This operation will unconfigure cluster HA. Cluster HA must be
configured on a two-node cluster to ensure data access availability in
the event of storage failover.
```

```
Do you want to continue? {y|n}: y
```

```
Notice: HA is disabled.
```

```
cluster_A::*>
```

```
cluster_B::*> cluster ha modify -configured false
```

```
Warning: This operation will unconfigure cluster HA. Cluster HA must be
configured on a two-node cluster to ensure data access availability in
the event of storage failover.
```

```
Do you want to continue? {y|n}: y
```

```
Notice: HA is disabled.
```

```
cluster_B::*>
```

4. Verify that cluster HA is disabled.



You must run this command on both clusters.

```
cluster_A::> cluster ha show
```

```
High Availability Configured: false
```

```
Warning: Cluster HA has not been configured. Cluster HA must be configured
```

```
on a two-node cluster to ensure data access availability in the event of storage failover. Use the "cluster ha modify -configured true" command to configure cluster HA.
```

```
cluster_A::>
```

```
cluster_B::> cluster ha show
```

```
High Availability Configured: false
```

```
Warning: Cluster HA has not been configured. Cluster HA must be configured
```

```
on a two-node cluster to ensure data access availability in the event of storage failover. Use the "cluster ha modify -configured true" command to configure cluster HA.
```

```
cluster_B::>
```

Joining the MetroCluster IP nodes to the clusters

You must add the four new MetroCluster IP nodes to the existing MetroCluster configuration.

About this task

You must perform this task on both clusters.

Steps

1. Add the MetroCluster IP nodes to the existing MetroCluster configuration.
 - a. Join the first MetroCluster IP node (node_A_3-IP) to the existing MetroCluster FC configuration.

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
```

```
"back" - if you want to change previously answered questions, and
```

```
"exit" or "quit" - if you want to quit the cluster setup wizard.
```

```
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster setup".
```

```
To accept a default or omit a question, do not enter a value.
```

```
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter autosupport modify -support
disable
within 24 hours.
```

```
Enabling AutoSupport can significantly speed problem determination
and
resolution, should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}: yes
```

```
Enter the node management interface port [e0M]:
Enter the node management interface IP address: 172.17.8.93
Enter the node management interface netmask: 255.255.254.0
Enter the node management interface default gateway: 172.17.8.1
A node management interface on port e0M with IP address 172.17.8.93
has been created.
```

```
Use your web browser to complete cluster setup by accessing
https://172.17.8.93
```

```
Otherwise, press Enter to complete cluster setup using the command
line
interface:
```

```
Do you want to create a new cluster or join an existing cluster?
{create, join}:
join
```

```
Existing cluster interface configuration found:
```

Port	MTU	IP	Netmask
e0c	9000	169.254.148.217	255.255.0.0
e0d	9000	169.254.144.238	255.255.0.0

```
Do you want to use this configuration? {yes, no} [yes]: yes
```

```
.
.
.
```

- b. Join the second MetroCluster IP node (node_A_4-IP) to the existing MetroCluster FC configuration.
2. Repeat these steps to join node_B_3-IP and node_B_4-IP to cluster_B.

3. If you are using the Onboard Key Manager, perform the following steps from the cluster where you've added a new node:
 - a. Synchronize the key manager configuration:

```
security key-manager onboard sync
```

- b. Enter the Onboard Key Manager passphrase when prompted.

Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates

You must create cluster peering LIFs, create the MetroCluster interfaces on the new MetroCluster IP nodes.

About this task

The home port used in the examples are platform-specific. You should use the appropriate home port specific to MetroCluster IP node platform.

Steps

1. On the new MetroCluster IP nodes, [configure the intercluster LIFs](#).
2. On each site, verify that cluster peering is configured:

```
cluster peer show
```

The following example shows the cluster peering configuration on cluster_A:

```
cluster_A:> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability
Authentication
-----
cluster_B                  1-80-000011          Available          ok
```

The following example shows the cluster peering configuration on cluster_B:

```
cluster_B:> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability
Authentication
-----
cluster_A 1-80-000011      Available          ok
```

3. Configure the DR group for the MetroCluster IP nodes:

```
metrocluster configuration-settings dr-group create -partner-cluster
```

```

cluster_A::> metrocluster configuration-settings dr-group create
-partner-cluster
cluster_B -local-node node_A_3-IP -remote-node node_B_3-IP
[Job 259] Job succeeded: DR Group Create is successful.
cluster_A::>

```

4. Verify that the DR group is created.

```
metrocluster configuration-settings dr-group show
```

```

cluster_A::> metrocluster configuration-settings dr-group show

DR Group ID Cluster          Node          DR Partner
Node
-----
2          cluster_A
          node_A_3-IP   node_B_3-IP
          node_A_4-IP   node_B_4-IP
          cluster_B
          node_B_3-IP   node_A_3-IP
          node_B_4-IP   node_A_4-IP

4 entries were displayed.

cluster_A::>

```

You will notice that the DR group for the old MetroCluster FC nodes (DR Group 1) is not listed when you run the `metrocluster configuration-settings dr-group show` command.

You can use `metrocluster node show` command on both sites to list all nodes.

```
cluster_A::> metrocluster node show
```

DR	Configuration	DR		
Group	Cluster	Node	State	Mirroring Mode
1	cluster_A			
		node_A_1-FC	configured	enabled normal
		node_A_2-FC	configured	enabled normal
	cluster_B			
		node_B_1-FC	configured	enabled normal
		node_B_2-FC	configured	enabled normal
2	cluster_A			
		node_A_3-IP	ready to configure	- -
		node_A_4-IP	ready to configure	- -

```
cluster_B::> metrocluster node show
```

DR	Configuration	DR		
Group	Cluster	Node	State	Mirroring Mode
1	cluster_B			
		node_B_1-FC	configured	enabled normal
		node_B_2-FC	configured	enabled normal
	cluster_A			
		node_A_1-FC	configured	enabled normal
		node_A_2-FC	configured	enabled normal
2	cluster_B			
		node_B_3-IP	ready to configure	- -
		node_B_4-IP	ready to configure	- -

5. Configure the MetroCluster IP interfaces for the newly joined MetroCluster IP nodes:



Do not use 169.254.17.x or 169.254.18.x IP addresses when you create MetroCluster IP interfaces to avoid conflicts with system auto-generated interface IP addresses in the same range.

```
metrocluster configuration-settings interface create -cluster-name
```

See [Configuring and connecting the MetroCluster IP interfaces](#) for considerations when configuring the IP interfaces.



You can configure the MetroCluster IP interfaces from either cluster.

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_3-IP -home-port ela -address
172.17.26.10 -netmask 255.255.255.0
```

```
[Job 260] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_3-IP -home-port elb -address
172.17.27.10 -netmask 255.255.255.0
```

```
[Job 261] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_4-IP -home-port ela -address
172.17.26.11 -netmask 255.255.255.0
```

```
[Job 262] Job succeeded: Interface Create is successful.
```

```
cluster_A::> :metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_4-IP -home-port elb -address
172.17.27.11 -netmask 255.255.255.0
```

```
[Job 263] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_3-IP -home-port ela -address
172.17.26.12 -netmask 255.255.255.0
```

```
[Job 264] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_3-IP -home-port elb -address
172.17.27.12 -netmask 255.255.255.0
```

```
[Job 265] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_4-IP -home-port ela -address
172.17.26.13 -netmask 255.255.255.0
```

```
[Job 266] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_4-IP -home-port elb -address
172.17.27.13 -netmask 255.255.255.0
```

```
[Job 267] Job succeeded: Interface Create is successful.
```

6. Verify the MetroCluster IP interfaces are created:

```
metrocluster configuration-settings interface show
```

```

cluster_A::>metrocluster configuration-settings interface show

DR
Config
Group Cluster Node      Network Address Netmask      Gateway
State
-----
-----
2      cluster_A
      node_A_3-IP
      Home Port: e1a
      172.17.26.10      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.10      255.255.255.0      -
completed
      node_A_4-IP
      Home Port: e1a
      172.17.26.11      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.11      255.255.255.0      -
completed
      cluster_B
      node_B_3-IP
      Home Port: e1a
      172.17.26.13      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.13      255.255.255.0      -
completed
      node_B_3-IP
      Home Port: e1a
      172.17.26.12      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.12      255.255.255.0      -
completed
8 entries were displayed.

cluster_A>

```

7. Connect the MetroCluster IP interfaces:

```
metrocluster configuration-settings connection connect
```



This command might take several minutes to complete.

```
cluster_A::> metrocluster configuration-settings connection connect
cluster_A::>
```

8. Verify the connections are properly established:

```
metrocluster configuration-settings connection show
```

```
cluster_A::> metrocluster configuration-settings connection show

DR          Source          Destination
Group Cluster Node    Network Address Network Address Partner Type
Config State
-----
2          cluster_A
          node_A_3-IP**
          Home Port: e1a
          172.17.26.10    172.17.26.11    HA Partner
completed
          Home Port: e1a
          172.17.26.10    172.17.26.12    DR Partner
completed
          Home Port: e1a
          172.17.26.10    172.17.26.13    DR Auxiliary
completed
          Home Port: e1b
          172.17.27.10    172.17.27.11    HA Partner
completed
          Home Port: e1b
          172.17.27.10    172.17.27.12    DR Partner
completed
          Home Port: e1b
          172.17.27.10    172.17.27.13    DR Auxiliary
completed
          node_A_4-IP
          Home Port: e1a
          172.17.26.11    172.17.26.10    HA Partner
completed
          Home Port: e1a
          172.17.26.11    172.17.26.13    DR Partner
completed
          Home Port: e1a
```

```

172.17.26.11    172.17.26.12    DR Auxiliary
completed
    Home Port: e1b
172.17.27.11    172.17.27.10    HA Partner
completed
    Home Port: e1b
172.17.27.11    172.17.27.13    DR Partner
completed
    Home Port: e1b
172.17.27.11    172.17.27.12    DR Auxiliary
completed

DR
Group Cluster Node    Source
Config State          Network Address
-----
-----
-----
-----
-----
2    cluster_B
    node_B_4-IP
    Home Port: e1a
172.17.26.13    172.17.26.12    HA Partner
completed
    Home Port: e1a
172.17.26.13    172.17.26.11    DR Partner
completed
    Home Port: e1a
172.17.26.13    172.17.26.10    DR Auxiliary
completed
    Home Port: e1b
172.17.27.13    172.17.27.12    HA Partner
completed
    Home Port: e1b
172.17.27.13    172.17.27.11    DR Partner
completed
    Home Port: e1b
172.17.27.13    172.17.27.10    DR Auxiliary
completed
    node_B_3-IP
    Home Port: e1a
172.17.26.12    172.17.26.13    HA Partner
completed
    Home Port: e1a
172.17.26.12    172.17.26.10    DR Partner
completed
    Home Port: e1a
172.17.26.12    172.17.26.11    DR Auxiliary

```

```
completed
      Home Port: elb
      172.17.27.12    172.17.27.13    HA Partner
completed
      Home Port: elb
      172.17.27.12    172.17.27.10    DR Partner
completed
      Home Port: elb
      172.17.27.12    172.17.27.11    DR Auxiliary
completed
24 entries were displayed.

cluster_A::>
```

9. Verify disk autoassignment and partitioning:

```
disk show -pool Pool1
```

```

cluster_A::> disk show -pool Pool1
          Usable          Disk      Container      Container
Disk      Size Shelf Bay Type      Type      Name
Owner
-----
-----
1.10.4          -      10   4 SAS      remote    -
node_B_2
1.10.13         -      10  13 SAS      remote    -
node_B_2
1.10.14         -      10  14 SAS      remote    -
node_B_1
1.10.15         -      10  15 SAS      remote    -
node_B_1
1.10.16         -      10  16 SAS      remote    -
node_B_1
1.10.18         -      10  18 SAS      remote    -
node_B_2
...
2.20.0      546.9GB      20   0 SAS      aggregate aggr0_rha1_a1
node_a_1
2.20.3      546.9GB      20   3 SAS      aggregate aggr0_rha1_a2
node_a_2
2.20.5      546.9GB      20   5 SAS      aggregate rha1_a1_aggr1
node_a_1
2.20.6      546.9GB      20   6 SAS      aggregate rha1_a1_aggr1
node_a_1
2.20.7      546.9GB      20   7 SAS      aggregate rha1_a2_aggr1
node_a_2
2.20.10     546.9GB      20  10 SAS      aggregate rha1_a1_aggr1
node_a_1
...
43 entries were displayed.
cluster_A::>

```



On systems configured for Advanced Drive Partitioning (ADP), the container type is "shared" rather than "remote" as shown in the example output.

10. Mirror the root aggregates:

```
storage aggregate mirror -aggregate aggr0_node_A_3_IP
```



You must complete this step on each MetroCluster IP node.

```

cluster_A::> aggr mirror -aggregate aggr0_node_A_3_IP

Info: Disks would be added to aggregate "aggr0_node_A_3_IP" on node
"node_A_3-IP"
      in the following manner:

      Second Plex

      RAID Group rg0, 3 disks (block checksum, raid_dp)

Physical                                                    Usable
Size      Position   Disk                Type                Size
-----
-----
-          dparity    4.20.0              SAS                  -
-          parity     4.20.3              SAS                  -
-          data       4.20.1              SAS                  546.9GB
558.9GB

Aggregate capacity available for volume use would be 467.6GB.

Do you want to continue? {y|n}: y

cluster_A::>

```

11. Verify that the root aggregates are mirrored:

```
storage aggregate show
```

```

cluster_A::> aggr show

Aggregate      Size Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
aggr0_node_A_1_FC
      349.0GB   16.84GB   95% online    1 node_A_1-FC
raid_dp,
mirrored,
normal

```

```

aggr0_node_A_2_FC
      349.0GB   16.84GB   95% online      1 node_A_2-FC
raid_dp,

mirrored,

normal
aggr0_node_A_3_IP
      467.6GB   22.63GB   95% online      1 node_A_3-IP
raid_dp,

mirrored,

normal
aggr0_node_A_4_IP
      467.6GB   22.62GB   95% online      1 node_A_4-IP
raid_dp,

mirrored,

normal
aggr_data_a1
      1.02TB    1.01TB    1% online       1 node_A_1-FC
raid_dp,

mirrored,

normal
aggr_data_a2
      1.02TB    1.01TB    1% online       1 node_A_2-FC
raid_dp,

mirrored,

```

Finalizing the addition of the MetroCluster IP nodes

You must incorporate the new DR group into the MetroCluster configuration and create mirrored data aggregates on the new nodes.

Steps

1. Configure the MetroCluster depending on whether there is a single or multiple data aggregates on both clusters:

If your MetroCluster configuration has...	Then do this...
---	-----------------

Multiple data aggregates on both clusters	<p>From any node's prompt, configure MetroCluster:</p> <pre>metrocluster configure <node-name></pre> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px;">  <p>You must run <code>metrocluster configure</code> and not <code>metrocluster configure -refresh true</code></p> </div>
A single mirrored data aggregate on both clusters	<p>a. From any node's prompt, change to the advanced privilege level:</p> <pre>set -privilege advanced</pre> <p>You must respond with <code>y</code> when you are prompted to continue into advanced mode and you see the advanced mode prompt (<code>*></code>).</p> <p>b. Configure the MetroCluster with the <code>-allow-with-one-aggregate true</code> parameter:</p> <pre>metrocluster configure -allow-with-one-aggregate true -node-name <node-name></pre> <p>c. Return to the admin privilege level:</p> <pre>set -privilege admin</pre>



The best practice is to have multiple mirrored data aggregates. When there is only one mirrored aggregate, there is less protection because the metadata volumes are located on the same aggregate rather than on separate aggregates.

2. Reboot each of the new nodes:

```
node reboot -node <node_name> -inhibit-takeover true
```



You don't need to reboot the nodes in a specific order, but you should wait until one node is fully booted and all connections are established before rebooting the next node.

3. Verify that the nodes are added to their DR group:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
```

DR	Configuration	DR
Group Cluster Node	State	Mirroring Mode

1	cluster_A	
	node-A-1-FC	configured enabled normal
	node-A-2-FC	configured enabled normal
	Cluster-B	
	node-B-1-FC	configured enabled normal
	node-B-2-FC	configured enabled normal
2	cluster_A	
	node-A-3-IP	configured enabled normal
	node-A-4-IP	configured enabled normal
	Cluster-B	
	node-B-3-IP	configured enabled normal
	node-B-4-IP	configured enabled normal

8 entries were displayed.

```
cluster_A::>
```

4. Create mirrored data aggregates on each of the new MetroCluster nodes:

```
storage aggregate create -aggregate aggregate-name -node node-name -diskcount  
no-of-disks -mirror true
```



You must create at least one mirrored data aggregate per site. It is recommended to have two mirrored data aggregates per site on MetroCluster IP nodes to host the MDV volumes, however a single aggregate per site is supported (but not recommended). It is acceptable that one site of the MetroCluster has a single mirrored data aggregate and the other site has more than one mirrored data aggregate.

The following example shows the creation of an aggregate on node_A_3-IP.

```
cluster_A::> storage aggregate create -aggregate data_a3 -node node_A_3-  
IP -diskcount 10 -mirror t
```

```
Info: The layout for aggregate "data_a3" on node "node_A_3-IP" would be:
```

```
First Plex
```

```
RAID Group rg0, 5 disks (block checksum, raid_dp)
```

```
Usable
```

```
Physical
```

```
Position
```

```
Disk
```

```
Type
```

```
Size
```

```

Size
-----
-----
-      dparity    5.10.15      SAS          -
-      parity     5.10.16      SAS          -
-      data       5.10.17      SAS          546.9GB
547.1GB
-      data       5.10.18      SAS          546.9GB
558.9GB
-      data       5.10.19      SAS          546.9GB
558.9GB

```

Second Plex

RAID Group rg0, 5 disks (block checksum, raid_dp)

```

Usable
Physical
Position  Disk          Type          Size
-----
-----
-      dparity    4.20.17      SAS          -
-      parity     4.20.14      SAS          -
-      data       4.20.18      SAS          546.9GB
547.1GB
-      data       4.20.19      SAS          546.9GB
547.1GB
-      data       4.20.16      SAS          546.9GB
547.1GB

```

Aggregate capacity available for volume use would be 1.37TB.

Do you want to continue? {y|n}: y

[Job 440] Job succeeded: DONE

cluster_A::>

5. Verify that all nodes in the cluster are healthy:

```
cluster show
```

The output should display `true` for the `health` field for all nodes.

6. Confirm that takeover is possible and the nodes are connected by running the following command on both clusters:

```
storage failover show
```

```
cluster_A::> storage failover show
```

Node	Partner	Takeover Possible	State Description
Node_FC_1	Node_FC_2	true	Connected to Node_FC_2
Node_FC_2	Node_FC_1	true	Connected to Node_FC_1
Node_IP_1	Node_IP_2	true	Connected to Node_IP_2
Node_IP_2	Node_IP_1	true	Connected to Node_IP_1

7. Confirm that all disks attached to the newly-joined MetroCluster IP nodes are present:

```
disk show
```

8. Verify the health of the MetroCluster configuration by running the following commands:

- a. `metrocluster check run`
- b. `metrocluster check show`
- c. `metrocluster interconnect mirror show`
- d. `metrocluster interconnect adapter show`

9. Move the MDV_CRS volumes from the old nodes to the new nodes in advanced privilege.

- a. Display the volumes to identify the MDV volumes:



If you have a single mirrored data aggregate per site then move both the MDV volumes to this single aggregate. If you have two or more mirrored data aggregates, then move each MDV volume to a different aggregate.

The following example shows the MDV volumes in the volume show output:

```

cluster_A::> volume show
Vserver    Volume                Aggregate    State    Type    Size
Available Used%
-----
...

cluster_A  MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_A
          aggr_b1            -          RW          -
- -
cluster_A  MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_B
          aggr_b2            -          RW          -
- -
cluster_A  MDV_CRS_d6b0b313ff5611e9837100a098544e51_A
          aggr_a1            online     RW          10GB
9.50GB    0%
cluster_A  MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
          aggr_a2            online     RW          10GB
9.50GB    0%
...
11 entries were displayed.mple

```

b. Set the advanced privilege level:

```
set -privilege advanced
```

c. Move the MDV volumes, one at a time:

```
volume move start -volume mdv-volume -destination-aggregate aggr-on-new-node
-vserver vserver-name
```

The following example shows the command and output for moving MDV_CRS_d6b0b313ff5611e9837100a098544e51_A to aggregate data_a3 on node_A_3.

```

cluster_A::*> vol move start -volume
MDV_CRS_d6b0b313ff5611e9837100a098544e51_A -destination-aggregate
data_a3 -vserver cluster_A

Warning: You are about to modify the system volume
        "MDV_CRS_d6b0b313ff5611e9837100a098544e51_A". This might
cause severe
        performance or stability problems. Do not proceed unless
directed to
        do so by support. Do you want to proceed? {y|n}: y
[Job 494] Job is queued: Move
"MDV_CRS_d6b0b313ff5611e9837100a098544e51_A" in Vserver "cluster_A"
to aggregate "data_a3". Use the "volume move show -vserver cluster_A
-volume MDV_CRS_d6b0b313ff5611e9837100a098544e51_A" command to view
the status of this operation.

```

- d. Use the volume show command to check that the MDV volume has been successfully moved:

```
volume show mdv-name
```

The following output shows that the MDV volume has been successfully moved.

```

cluster_A::*> vol show MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
Vserver      Volume      Aggregate    State      Type      Size
Available Used%
-----
-----
cluster_A    MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
              aggr_a2      online      RW         10GB
9.50GB      0%

```

- e. Return to admin mode:

```
set -privilege admin
```

Moving the data to the new drive shelves

During the transition, you move data from the drive shelves in the MetroCluster FC configuration to the new MetroCluster IP configuration.

Before you begin

You should create new SAN LIFs on the destination or IP nodes and connect hosts prior to moving volumes to new the new aggregates.

1. To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.

- a. Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=end`
- b. Repeat the command on the partner cluster.

2. Move the data volumes to aggregates on the new controllers, one volume at a time.

Use the procedure in [Creating an aggregate and moving volumes to the new nodes](#).

3. Create SAN LIFs on the recently added nodes.

Use the following procedure in [Updating LUN paths for the new nodes](#).

4. Check if there are any node locked licenses on the FC nodes, if there are, they need to be added to the newly added nodes.

Use the following procedure in [Adding node-locked licenses](#).

5. Migrate the data LIFs.

Use the procedure in [Moving non-SAN data LIFs and cluster management LIFs to the new nodes](#) but do **not** perform the last two steps to migrate cluster management LIFs.



- You cannot migrate a LIF that is used for copy-offload operations with VMware vStorage APIs for Array Integration (VAAI).
- After you complete the transition of your MetroCluster nodes from FC to IP, you might need to move your iSCSI host connections to the new nodes, see [Moving Linux iSCSI hosts from MetroCluster FC to MetroCluster IP nodes](#).

Removing the MetroCluster FC controllers

You must perform clean-up tasks and remove the old controller modules from the MetroCluster configuration.

1. To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway.

- a. Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours`

`maintenance-window-in-hours` specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period: `system node autosupport invoke -node * -type all -message MAINT=end`

- b. Repeat the command on the partner cluster.

2. Identify the aggregates hosted on the MetroCluster FC configuration that need to be deleted.

In this example the following data aggregates are hosted by the MetroCluster FC cluster_B and need to be deleted: `aggr_data_a1` and `aggr_data_a2`.



You need to perform the steps to identify, offline and delete the data aggregates on both the clusters. The example is for one cluster only.

```
cluster_B::> aggr show
```

Aggregate Status	Size	Available	Used%	State	#Vols	Nodes	RAID
-----	-----	-----	-----	-----	-----	-----	-----
aggr0_node_A_1-FC	349.0GB	16.83GB	95%	online	1	node_A_1-FC	
raid_dp,							
mirrored,							
normal							
aggr0_node_A_2-FC	349.0GB	16.83GB	95%	online	1	node_A_2-FC	
raid_dp,							
mirrored,							
normal							
aggr0_node_A_3-IP	467.6GB	22.63GB	95%	online	1	node_A_3-IP	
raid_dp,							
mirrored,							
normal							
aggr0_node_A_3-IP	467.6GB	22.62GB	95%	online	1	node_A_4-IP	
raid_dp,							
mirrored,							
normal							
aggr_data_a1	1.02TB	1.02TB	0%	online	0	node_A_1-FC	
raid_dp,							
mirrored,							
normal							
aggr_data_a2	1.02TB	1.02TB	0%	online	0	node_A_2-FC	
raid_dp,							

```

mirrored,

normal
aggr_data_a3
      1.37TB      1.35TB      1% online      3 node_A_3-IP
raid_dp,

mirrored,

normal
aggr_data_a4
      1.25TB      1.24TB      1% online      2 node_A_4-IP
raid_dp,

mirrored,

normal
8 entries were displayed.

```

```
cluster_B::>
```

3. Check if the data aggregates on the FC nodes have any MDV_aud volumes, and delete them prior to deleting the aggregates.

You must delete the MDV_aud volumes as they cannot be moved.

4. Take each of the data aggregates offline, and then delete them:

- a. Take the aggregate offline: `storage aggregate offline -aggregate aggregate-name`

The following example shows the aggregate `aggr_data_a1` being taken offline:

```

cluster_B::> storage aggregate offline -aggregate aggr_data_a1

Aggregate offline successful on aggregate: aggr_data_a1

```

- b. Delete the aggregate: `storage aggregate delete -aggregate aggregate-name`

You can destroy the plex when prompted.

The following example shows the aggregate `aggr_data_a1` being deleted.

```

cluster_B::> storage aggregate delete -aggregate aggr_data_a1
Warning: Are you sure you want to destroy aggregate "aggr_data_a1"?
{y|n}: y
[Job 123] Job succeeded: DONE

cluster_B::>

```

5. Identify the MetroCluster FC DR group that need to be removed.

In the following example the MetroCluster FC nodes are in DR Group '1', and this is the DR group that need to be removed.

```

cluster_B::> metrocluster node show

DR
Group Cluster Node          Configuration State      DR
Mirroring Mode
-----
1      cluster_A
      node_A_1-FC             configured enabled   normal
      node_A_2-FC             configured enabled   normal
      cluster_B
      node_B_1-FC             configured enabled   normal
      node_B_2-FC             configured enabled   normal
2      cluster_A
      node_A_3-IP             configured enabled   normal
      node_A_4-IP             configured enabled   normal
      cluster_B
      node_B_3-IP             configured enabled   normal
      node_B_3-IP             configured enabled   normal
8 entries were displayed.

cluster_B::>

```

6. Move the cluster management LIF from a MetroCluster FC node to a MetroCluster IP node:

```

cluster_B::> network interface migrate -vserver svm-name -lif cluster_mgmt
-destination-node node-in-metrocluster-ip-dr-group -destination-port
available-port

```

7. Change the home node and home port of the cluster management LIF: `cluster_B::> network interface modify -vserver svm-name -lif cluster_mgmt -service-policy default-management -home-node node-in-metrocluster-ip-dr-group -home-port lif-port`

8. Move epsilon from a MetroCluster FC node to a MetroCluster IP node:

- a. Identify which node currently has epsilon: `cluster show -fields epsilon`

```

cluster_B::> cluster show -fields epsilon
node          epsilon
-----
node_A_1-FC   true
node_A_2-FC   false
node_A_1-IP   false
node_A_2-IP   false
4 entries were displayed.

```

- b. Set epsilon to false on the MetroCluster FC node (node_A_1-FC): `cluster modify -node fc-node -epsilon false`
- c. Set epsilon to true on the MetroCluster IP node (node_A_1-IP): `cluster modify -node ip-node -epsilon true`
- d. Verify that epsilon has moved to the correct node: `cluster show -fields epsilon`

```

cluster_B::> cluster show -fields epsilon
node          epsilon
-----
node_A_1-FC   false
node_A_2-FC   false
node_A_1-IP   true
node_A_2-IP   false
4 entries were displayed.

```

9. Modify the IP address for the cluster peer of the transitioned IP nodes for each cluster:

- a. Identify the cluster_A peer by using the `cluster peer show` command:

```

cluster_A::> cluster peer show
Peer Cluster Name      Cluster Serial Number Availability
Authentication
-----
cluster_B              1-80-000011          Unavailable      absent

```

- i. Modify the cluster_A peer IP address:

```

cluster peer modify -cluster cluster_A -peer-addr node_A_3_IP -address
-family ipv4

```

- b. Identify the cluster_B peer by using the `cluster peer show` command:

```

cluster_B::> cluster peer show
Peer Cluster Name          Cluster Serial Number Availability
Authentication
-----
cluster_A                  1-80-000011          Unavailable    absent

```

i. Modify the cluster_B peer IP address:

```

cluster peer modify -cluster cluster_B -peer-addr node_B_3_IP -address
-family ipv4

```

c. Verify that the cluster peer IP address is updated for each cluster:

i. Verify that the IP address is updated for each cluster by using the `cluster peer show -instance` command.

The Remote Intercluster Addresses field in the following examples displays the updated IP address.

Example for cluster_A:

```

cluster_A::> cluster peer show -instance

Peer Cluster Name: cluster_B
      Remote Intercluster Addresses: 172.21.178.204,
172.21.178.212
      Availability of the Remote Cluster: Available
      Remote Cluster Name: cluster_B
      Active IP Addresses: 172.21.178.212,
172.21.178.204
      Cluster Serial Number: 1-80-000011
      Remote Cluster Nodes: node_B_3-IP,
node_B_4-IP
      Remote Cluster Health: true
      Unreachable Local Nodes: -
      Address Family of Relationship: ipv4
      Authentication Status Administrative: use-authentication
      Authentication Status Operational: ok
      Last Update Time: 4/20/2023 18:23:53
      IPspace for the Relationship: Default
      Proposed Setting for Encryption of Inter-Cluster Communication: -
      Encryption Protocol For Inter-Cluster Communication: tls-psk
      Algorithm By Which the PSK Was Derived: jpake

cluster_A::>

```

Example for cluster_B

```
cluster_B::> cluster peer show -instance

                Peer Cluster Name: cluster_A
    Remote Intercluster Addresses: 172.21.178.188,
172.21.178.196 <<<<<<< Should reflect the modified address
    Availability of the Remote Cluster: Available
                Remote Cluster Name: cluster_A
                Active IP Addresses: 172.21.178.196,
172.21.178.188
                Cluster Serial Number: 1-80-000011
                Remote Cluster Nodes: node_A_3-IP,
                                     node_A_4-IP
                Remote Cluster Health: true
                Unreachable Local Nodes: -
                Address Family of Relationship: ipv4
    Authentication Status Administrative: use-authentication
    Authentication Status Operational: ok
                Last Update Time: 4/20/2023 18:23:53
                IPspace for the Relationship: Default
    Proposed Setting for Encryption of Inter-Cluster Communication: -
    Encryption Protocol For Inter-Cluster Communication: tls-psk
    Algorithm By Which the PSK Was Derived: jpake

cluster_B::>
```

10. On each cluster, remove the DR group containing the old nodes from the MetroCluster FC configuration.

You must perform this step on both clusters, one at a time.

```
cluster_B::> metrocluster remove-dr-group -dr-group-id 1
```

Warning: Nodes in the DR group that are removed from the MetroCluster configuration will lose their disaster recovery protection.

Local nodes "node_A_1-FC, node_A_2-FC" will be removed from the MetroCluster configuration. You must repeat the operation on the partner cluster "cluster_B" to remove the remote nodes in the DR group.

Do you want to continue? {y|n}: y

Info: The following preparation steps must be completed on the local and partner clusters before removing a DR group.

1. Move all data volumes to another DR group.
2. Move all MDV_CRS metadata volumes to another DR group.
3. Delete all MDV_aud metadata volumes that may exist in the DR group to be removed.
4. Delete all data aggregates in the DR group to be removed. Root aggregates are not deleted.
5. Migrate all data LIFs to home nodes in another DR group.
6. Migrate the cluster management LIF to a home node in another DR group. Node management and inter-cluster LIFs are not migrated.
7. Transfer epsilon to a node in another DR group.

The command is vetoed if the preparation steps are not completed on the local and partner clusters.

Do you want to continue? {y|n}: y

[Job 513] Job succeeded: Remove DR Group is successful.

```
cluster_B::>
```

11. Verify that the nodes are ready to be removed from the clusters.

You must perform this step on both clusters.



At this point, the `metrocluster node show` command only shows the local MetroCluster FC nodes and no longer shows the nodes that are part of the partner cluster.

```
cluster_B::> metrocluster node show
```

```
DR
Group Cluster Node Configuration State DR
-----
-----
-----
1 cluster_A
  node_A_1-FC ready to configure - -
  node_A_2-FC ready to configure - -
2 cluster_A
  node_A_3-IP configured enabled normal
  node_A_4-IP configured enabled normal
  cluster_B
  node_B_3-IP configured enabled normal
  node_B_4-IP configured enabled normal
6 entries were displayed.

cluster_B::>
```

12. Disable storage failover for the MetroCluster FC nodes.

You must perform this step on each node.

```
cluster_A::> storage failover modify -node node_A_1-FC -enabled false
cluster_A::> storage failover modify -node node_A_2-FC -enabled false
cluster_A::>
```

13. Unjoin the MetroCluster FC nodes from the clusters: `cluster unjoin -node node-name`

You must perform this step on each node.

```
cluster_A::> cluster unjoin -node node_A_1-FC
```

```
Warning: This command will remove node "node_A_1-FC" from the cluster.  
You must
```

```
    remove the failover partner as well. After the node is removed,  
erase
```

```
    its configuration and initialize all disks by using the "Clean  
    configuration and initialize all disks (4)" option from the  
boot menu.
```

```
Do you want to continue? {y|n}: y
```

```
[Job 553] Job is queued: Cluster remove-node of Node:node_A_1-FC with  
UUID:6c87de7e-ff54-11e9-8371
```

```
[Job 553] Checking prerequisites
```

```
[Job 553] Cleaning cluster database
```

```
[Job 553] Job succeeded: Node remove succeeded
```

```
If applicable, also remove the node's HA partner, and then clean its  
configuration and initialize all disks with the boot menu.
```

```
Run "debug vreport show" to address remaining aggregate or volume  
issues.
```

```
cluster_B::>
```

14. If the configuration uses FC-to-SAS bridges or FC back-end switches, disconnect and remove them.

Remove FC-to-SAS bridges

- a. Identify the bridges:

```
system bridge show
```

- b. Remove the bridges:

```
system bridge remove -name <bridge_name>
```

- c. Confirm the bridges are removed:

```
system bridge show
```

The following example shows that the bridges are removed:

Example

```
cluster1::> system bridge remove -name ATTO_10.226.197.16
cluster1::> system bridge show

Is          Monitor
  Bridge    Symbolic Name Vendor  Model      Bridge WWN
Monitored Status
-----
ATTO_FibreBridge6500N_1
          Bridge Number 16
                   Atto   FibreBridge 6500N
                                2000001086603824
false     -
          ATTO_FibreBridge6500N_2
                   Not Set   Atto   FibreBridge 6500N
                                20000010866037e8
false     -
          ATTO_FibreBridge6500N_3
                   Not Set   Atto   FibreBridge 6500N
                                2000001086609e0e
false     -
          ATTO_FibreBridge6500N_4
                   Not Set   Atto   FibreBridge 6500N
                                2000001086609c06
false     -
          4 entries were displayed.
```

Remove FC switches

a. Identify the switches:

```
system switch fibre-channel show
```

b. Remove the switches:

```
system switch fibre-channel remove -switch-name <switch_name>
```

c. Confirm the switches are removed:

```
system switch fibre-channel show
```

Example

```
cluster1::> system switch fibre-channel show
      Symbolic                               Is
Monitor
  Switch      Name      Vendor  Model      Switch WWN
Monitored Status
-----
Cisco_10.226.197.34
      mcc-cisco-8Gb-fab-4
      Cisco    DS-C9148-16P-K9
      2000547fee78f088
true      ok
      mcc-cisco-8Gb-fab-1
      mcc-cisco-8Gb-fab-1
      Cisco    -
false     -
      mcc-cisco-8Gb-fab-2
      mcc-cisco-8Gb-fab-2
      Cisco    -
false     -
      mcc-cisco-8Gb-fab-3
      mcc-cisco-8Gb-fab-3
      Cisco    -
false     -
      4 entries were displayed.
cluster1::> system switch fibre-channel remove -switch-name
Cisco_10.226.197.34
cluster1::> system switch fibre-channel show
      Symbolic                               Is
Monitor
  Switch      Name      Vendor  Model      Switch WWN
Monitored Status
-----
      mcc-cisco-8Gb-fab-4
      mcc-cisco-8Gb-fab-4
      Cisco
      -
false     -
      mcc-cisco-8Gb-fab-1
      mcc-cisco-8Gb-fab-1
      Cisco    -
false     -
      mcc-cisco-8Gb-fab-2
```

```

                mcc-cisco-8Gb-fab-2
                    Cisco  -      -
false  -
    mcc-cisco-8Gb-fab-3
                mcc-cisco-8Gb-fab-3
                    Cisco  -      -
false  -
    4 entries were displayed
cluster1::>

```

15. Power down the MetroCluster FC controller modules and storage shelves.
16. Disconnect and remove the MetroCluster FC controller modules and storage shelves.

Completing the transition

To complete the transition you must verify the operation of the new MetroCluster IP configuration.

1. Verify the MetroCluster IP configuration.

You must perform this step on each cluster in advanced privilege mode.

The following example shows the output for cluster_A.

```

cluster_A::> cluster show
Node                Health  Eligibility  Epsilon
-----
node_A_1-IP        true   true         false
node_A_2-IP        true   true         false
2 entries were displayed.

cluster_A::>

```

The following example shows the output for cluster_B.

```

cluster_B::> cluster show
Node                Health  Eligibility  Epsilon
-----
node_B_1-IP        true   true         false
node_B_2-IP        true   true         false
2 entries were displayed.

cluster_B::>

```

2. Enable cluster HA and storage failover.

You must perform this step on each cluster.

3. Verify that cluster HA capability is enabled.

```
cluster_A::> cluster ha show
High Availability Configured: true

cluster_A::>

cluster_A::> storage failover show

Node           Partner           Takeover
-----
node_A_1-IP    node_A_2-IP      true    Connected to node_A_2-IP
node_A_2-IP    node_A_1-IP      true    Connected to node_A_1-IP
2 entries were displayed.

cluster_A::>
```

4. Disable MetroCluster transition mode.

a. Change to the advanced privilege level: `set -privilege advanced`

b. Disable transition mode: `metrocluster transition disable`

c. Return to the admin privilege level: `set -privilege admin`

```
cluster_A::*> metrocluster transition disable

cluster_A::*>
```

5. Verify that transition is disabled: `metrocluster transition show-mode`

You must perform these steps on both clusters.

```
cluster_A::> metrocluster transition show-mode
Transition Mode
-----
not-enabled

cluster_A::>
```

```

cluster_B::> metrocluster transition show-mode
Transition Mode
-----
not-enabled

cluster_B::>

```

6. If you have an eight-node configuration, you must repeat the entire procedure starting from [Prepare for transition from a MetroCluster FC to a MetroCluster IP configuration](#) for each of the FC DR groups.

Sending a custom AutoSupport message after maintenance

After completing the transition, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

1. To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.
 - a. Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=end`
 - b. Repeat the command on the partner cluster.

Restoring Tiebreaker or Mediator monitoring

After completing the transition of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

1. Use the appropriate procedure for your configuration.

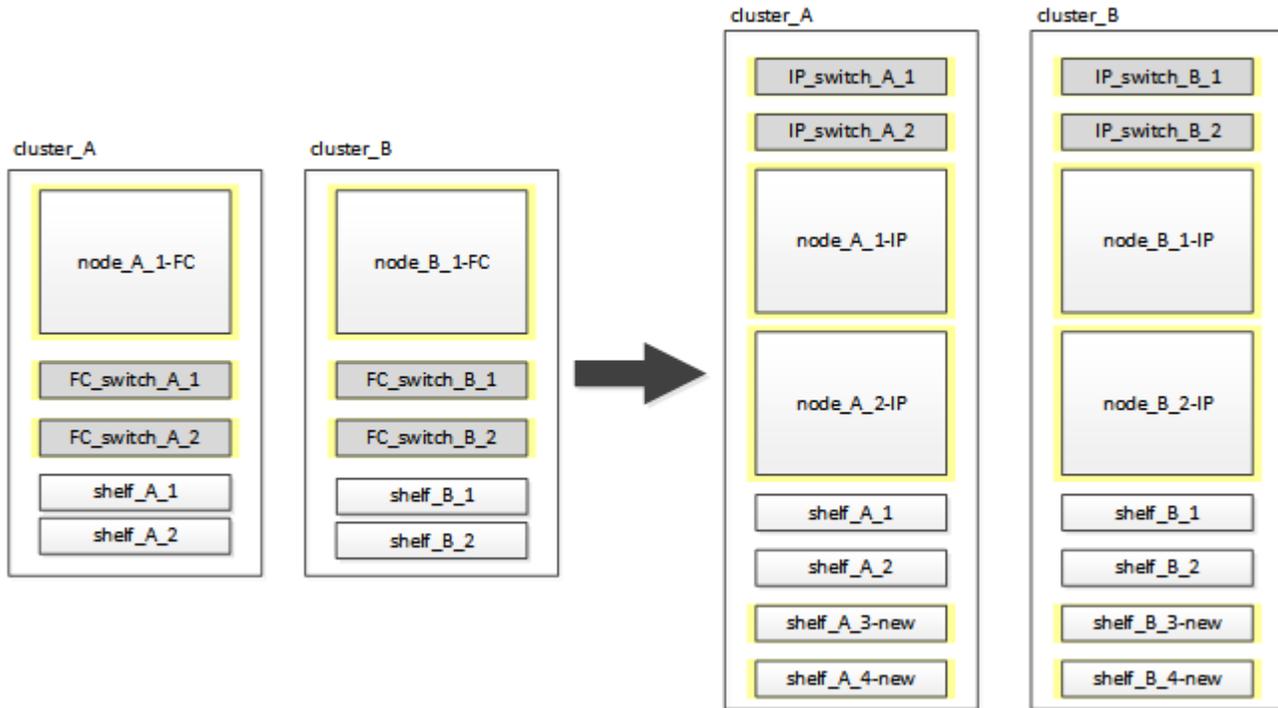
If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations
Mediator	Configure ONTAP Mediator from a MetroCluster IP configuration

Disruptively transition from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Disruptively transitioning from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Beginning with ONTAP 9.8, you can transition workloads and data from an existing two-node MetroCluster FC configuration to a new four-node MetroCluster IP configuration. Disk shelves from the MetroCluster FC nodes are moved to the IP nodes.

The following illustration provides a simplified view of the configuration before and after this transition procedure.



- This procedure is supported on systems running ONTAP 9.8 and later.
- This procedure is disruptive.
- This procedure applies only to a two-node MetroCluster FC configuration.

If you have a four-node MetroCluster FC configuration, see [Choosing your transition procedure](#).

- ADP is not supported on the four-node MetroCluster IP configuration created by this procedure.
- You must meet all requirements and follow all steps in the procedure.
- The existing storage shelves are moved to the new MetroCluster IP nodes.
- Additional storage shelves can be added to the configuration if necessary.

See [Drive shelf reuse and drive requirements for disruptive FC-to-IP transition](#).

Example naming in this procedure

This procedure uses example names throughout to identify the DR groups, nodes, and switches involved.

The nodes in the original configuration have the suffix -FC, indicating that they are in a fabric-attached or stretch MetroCluster configuration.

Components	cluster_A at site_A	cluster_B at site_B
------------	---------------------	---------------------

dr_group_1-FC	<ul style="list-style-type: none"> • node_A_1-FC • shelf_A_1 • shelf_A_2 	<ul style="list-style-type: none"> • node_B_1-FC • shelf_B_1 • shelf_B_2
dr_group_2-IP	<ul style="list-style-type: none"> • node_A_1-IP • node_A_2-IP • shelf_A_1 • shelf_A_2 • shelf_A_3-new • shelf_A_4-new 	<ul style="list-style-type: none"> • node_B_1-IP • node_B_2-IP • shelf_B_1 • shelf_B_2 • shelf_B_3-new • shelf_B_4-new
Switches	<ul style="list-style-type: none"> • switch_A_1-FC • switch_A_2-FC • switch_A_1-IP • switch_A_2-IP 	<ul style="list-style-type: none"> • switch_B_1-FC • switch_B_2-FC • switch_B_1-IP • switch_B_2-IP

Preparing for disruptive FC-to-IP transition

Before starting the transition process, you must make sure the configuration meets the requirements.

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

General requirements for disruptive FC-to-IP transition

The existing MetroCluster FC configuration must meet the following requirements:

- It must be a two-node configuration and all nodes must be running ONTAP 9.8 or later.

It can be a two-node fabric-attached or stretched MetroCluster.

- It must meet all requirements and cabling as described in the *MetroCluster Installation and Configuration* procedures.

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

- It cannot be configured with NetApp Storage Encryption (NSE).
- The MDV volumes cannot be encrypted.

You must have remote console access for all six nodes from either MetroCluster site or plan for travel between the sites as required by the procedure.

Drive shelf reuse and drive requirements for disruptive FC-to-IP transition

You must ensure that adequate spare drives and root aggregate space is available on the storage shelves.

Reusing the existing storage shelves

When using this procedure, the existing storage shelves are retained for use by the new configuration. When node_A_1-FC and node_B_1-FC are removed, the existing drive shelves are connected to node_A_1-IP and node_A_2-IP on cluster_A and to node_B_1-IP and node_B_2-IP on cluster_B.

- The existing storage shelves (those attached to node_A_1-FC and node_B_1-FC) must be supported by the new platform models.

If the existing shelves are not supported by the new platform models, see [Disruptively transitioning when existing shelves are not supported on new controllers \(ONTAP 9.8 and later\)](#).

- You must ensure you don't exceed the platform limits for drives, etc.

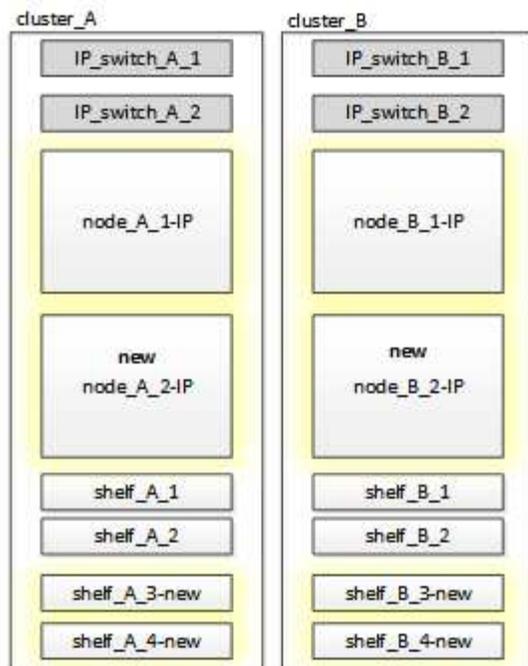
[NetApp Hardware Universe](#)

Storage requirements for the additional controllers

Additional storage must be added, if necessary, to accommodate the two additional controllers (node_A_2-IP and node_B_2-ip), because the configuration is changing from a two-node to a four-node arrangement.

- Depending on the spare drives available in the existing shelves, additional drives must be added to accommodate the additional controllers in the configuration.

This might require additional storage shelves, as shown in the following illustration.



You need to have additional 14 - 18 drives each for the third and fourth controllers (node_A_2-IP and node_B_2-IP):

- Three pool0 drives
- Three pool1 drives
- Two spare drives
- Six to ten drives for the system volume
- You must ensure that the configuration, including the new nodes, does not exceed the platform limits for the configuration, including drive count, root aggregate size capacity, etc.

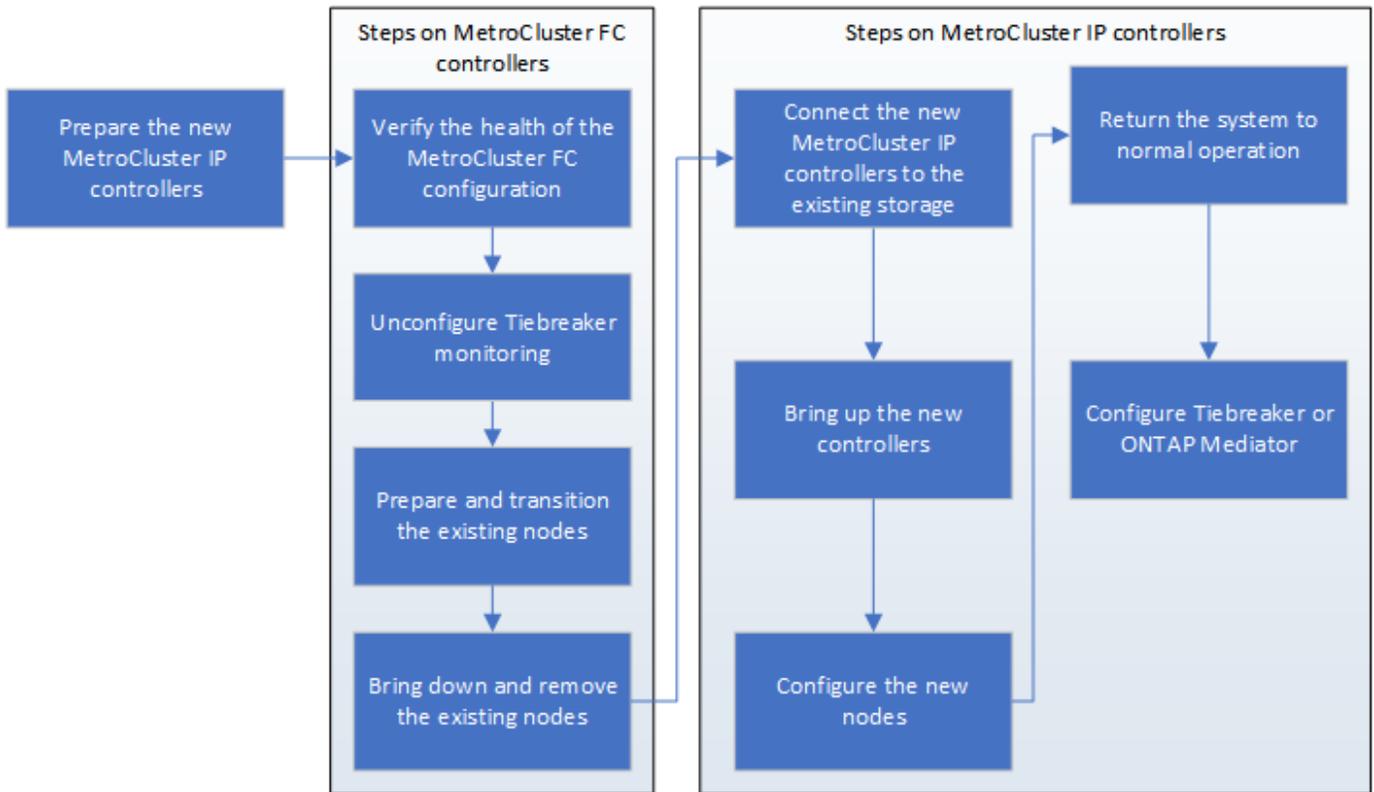
This information is available for each platform model at *NetApp Hardware Universe*.

[NetApp Hardware Universe](#)

Workflow for disruptive transition

You must follow the specific workflow to ensure a successful transition.

As you prepare for the transition, plan for travel between the sites. Note that after the remote nodes are racked and cabled, you need serial terminal access to the nodes. Service Processor access is not be available until the nodes are configured.



Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes

You must adjust the port and LIF configuration of the MetroCluster FC node so it is compatible with that of the MetroCluster IP node that will replace it.

About this task

When the new nodes are first booted during the upgrade process, each node uses the most recent configuration of the node it is replacing. When you boot node_A_1-IP, ONTAP attempts to host LIFs on the same ports that were used on node_A_1-FC.

During the transition procedure, you will perform steps on both the old and new nodes to ensure correct cluster, management, and data LIF configuration.

Steps

1. Identify any conflicts between the existing MetroCluster FC port usage and the port usage for the MetroCluster IP interfaces on the new nodes.

You must identify the MetroCluster IP ports on the new MetroCluster IP controllers using the table below. Then check and record if any data LIFs or cluster LIFs exist on those ports on the MetroCluster FC nodes.

These conflicting data LIFs or cluster LIFs on the MetroCluster FC nodes will be moved at the appropriate step in the transition procedure.

The following table shows the MetroCluster IP ports by platform model. You can ignore the VLAN ID column.

Platform model	MetroCluster IP port	VLAN ID	
----------------	----------------------	---------	--

AFF A800	e0b	Not used	
	e1b		
AFF A700 and FAS9000	e5a		
	e5b		
AFF A320	e0g		
	e0h		
AFF A300 and FAS8200	e1a		
	e1b		
FAS8300/A400/FAS8700	e1a		10
	e1b		20
AFF A250 and FAS500f	e0c	10	
	e0b	20	

You can fill in the following table and refer to it later in the transition procedure.

Ports	Corresponding MetroCluster IP interface ports (from table above)	Conflicting LIFs on these ports on the MetroCluster FC nodes
First MetroCluster IP port on node_A_1-FC		
Second MetroCluster IP port on node_A_1-FC		
First MetroCluster IP port on node_B_1-FC		
Second MetroCluster IP port on node_B_1-FC		

- Determine which physical ports are available on the new controllers and which LIFs can be hosted on the ports.

The controller's port usage depends on the platform model and IP switch model you will use in the MetroCluster IP configuration. You can gather the port usage of the new platforms from the *NetApp*

NetApp Hardware Universe

3. If desired, record the port information for node_A_1-FC and node_A_1-IP.

You will refer to the table as you carry out the transition procedure.

In the columns for node_A_1-IP, add the physical ports for the new controller module and plan the IPspaces and broadcast domains for the new node.

LIF	node_A_1-FC			node_A_1-IP		
	Ports	IPspaces	Broadcast domains	Ports	IPspaces	Broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

4. If desired, record all the port information for node_B_1-FC.

You will refer to the table as you carry out the upgrade procedure.

In the columns for node_B_1-IP, add the physical ports for the new controller module and plan the LIF port usage, IPspaces and broadcast domains for the new node.

	node_B_1-FC			node_B_1-IP		
LIF	Physical ports	IPspaces	Broadcast domains	Physical ports	IPspaces	Broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

Preparing the MetroCluster IP controllers

You must prepare the four new MetroCluster IP nodes and install the correct ONTAP version.

About this task

This task must be performed on each of the new nodes:

- node_A_1-IP
- node_A_2-IP
- node_B_1-IP
- node_B_2-IP

The nodes should be connected to any **new** storage shelves. They must **not** be connected to the existing storage shelves containing data.

These steps can be performed now, or later in the procedure when the controllers and shelves are racked. In

any case, you must make sure you clear the configuration and prepare the nodes **before** connecting them to the existing storage shelves and **before** making any configuration changes to the MetroCluster FC nodes.



Do not perform these steps with the MetroCluster IP controllers connected to the existing storage shelves that were connected to the MetroCluster FC controllers.

In these steps, you clear the configuration on the nodes and clear the mailbox region on new drives.

Steps

1. Connect the controller modules to the new storage shelves.
2. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be “mccip”.

3. If the displayed system state of the controller or chassis is not correct, set the HA state:

```
ha-config modify controller mccip`ha-config modify chassis mccip
```

4. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

5. Repeat the following substeps on all four nodes to clear the configuration:
 - a. Set the environmental variables to default values:

```
set-defaults
```

- b. Save the environment:

```
saveenv
```

```
bye
```

6. Repeat the following substeps to boot all four nodes using the 9a option on the boot menu.
 - a. At the LOADER prompt, launch the boot menu:

```
boot_ontap menu
```

- b. At the boot menu, select option “9a” to reboot the controller.

7. Boot each of the four nodes to Maintenance mode using option “5” on the boot menu.
8. Record the system ID and from each of the four nodes:

```
sysconfig
```

9. Repeat the following steps on node_A_1-IP and node_B_1-IP.
 - a. Assign ownership of all disks local to each site:

```
disk assign adapter.xx.*
```

- b. Repeat the previous step for each HBA with attached drive shelves on node_A_1-IP and node_B_1-IP.
10. Repeat the following steps on node_A_1-IP and node_B_1-IP to clear the mailbox region on each local disk.
 - a. Destroy the mailbox region on each disk:

```
mailbox destroy local``mailbox destroy partner
```

11. Halt all four controllers:

```
halt
```

12. On each controller, display the boot menu:

```
boot_ontap menu
```

13. On each of the four controllers, clear the configuration:

```
wipeconfig
```

When the wipeconfig operation completes, the node automatically returns to the boot menu.

14. Repeat the following substeps to again boot all four nodes using the 9a option on the boot menu.

- a. At the LOADER prompt, launch the boot menu:

```
boot_ontap menu
```

- b. At the boot menu, select option "9a" to reboot the controller.
- c. Let the controller module complete booting before moving to the next controller module.

After "9a" completes, the nodes automatically return to the boot menu.

15. Power off the controllers.

Verifying the health of the MetroCluster FC configuration

You must verify the health and connectivity of the MetroCluster FC configuration prior to performing the transition

This task is performed on the MetroCluster FC configuration.

1. Verify the operation of the MetroCluster configuration in ONTAP:
 - a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

2. Verify that the nodes are in non-HA mode:

```
storage failover show
```

Removing the existing configuration from the Tiebreaker or other monitoring software

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the Tiebreaker or other software prior to transition.

Steps

1. Remove the existing MetroCluster configuration from the Tiebreaker software.

[Removing MetroCluster configurations](#)

2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

Transitioning the MetroCluster FC nodes

You must gather information from the existing MetroCluster FC nodes, send an autosupport message announcing the start of maintenance, and transition the nodes.

Gathering information from the existing controller modules before the transition

Before transitioning, you must gather information for each of the nodes.

This task is performed on the existing nodes:

- node_A_1-FC
- node_B_1-FC

1. Gather the output for the commands in the following table.

Category	Commands	Notes
License	system license show	
Shelves and numbers of disks in each shelf and flash storage details and memory and NVRAM and and network cards	system node run -node node_name sysconfig	
Cluster network and node management LIFs	system node run -node node_name sysconfig network interface show -role "cluster,node-mgmt,data"	
SVM information	vserver show	
Protocol information	nfs show iscsi show cifs show	
Physical ports	network port show -node node_name -type physical network port show	
Failover Groups	network interface failover-groups show -vserver vserver_name	Record the names and ports of failover groups that are not clusterwide.
VLAN configuration	network port vlan show -node node_name	Record each network port and VLAN ID pairing.
Interface group configuration	network port ifgrp show -node node_name -instance	Record the names of the interface groups and the ports assigned to them.
Broadcast domains	network port broadcast-domain show	
IPspace	network ipspace show	
Volume info	volume show and volume show -fields encrypt	
Aggregate Info	storage aggregate show and storage aggr encryption show andstorage aggregate object-store show	
Disk ownership information	storage aggregate show and storage aggr encryption show andstorage aggregate object-store show	
Encryption	storage failover mailbox-disk show and security key-manager backup show	Also preserve the passphrase used to enable key-manager. In the case of external key-manager you will need the authentication information for the client and server.

Category	Commands	Notes
Encryption	security key-manager show	
Encryption	security key-manager external show	
Encryption	systemshell local kenv kmip.init.ipaddr ip-address	
Encryption	systemshell local kenv kmip.init.netmask netmask	
Encryption	systemshell local kenv kmip.init.gateway gateway	
Encryption	systemshell local kenv kmip.init.interface interface	

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. This prevents them from opening a case on the assumption that a disruption has occurred.

This task must be performed on each MetroCluster site.

1. To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway.
 - a. Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours`

`maintenance-window-in-hours` specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:
`system node autosupport invoke -node * -type all -message MAINT=end`
 - b. Repeat the command on the partner cluster.

Transitioning, shutting down, and removing the MetroCluster FC nodes

In addition to issuing commands on the MetroCluster FC nodes, this task includes physical uncabbling and removal of the controller modules at each site.

This task must be performed on each of the old nodes:

- node_A_1-FC
- node_B_1-FC

Steps

1. Stop all client traffic.
2. On either of the MetroCluster FC nodes, for example node_A_1-FC, enable transition.
 - a. Set the advanced privilege level: `set -priv advanced`

b. Enable transition: `metrocluster transition enable -transition-mode disruptive`

c. Return to admin mode: `set -priv admin`

3. Unmirror the root aggregate by deleting the remote plex of the root aggregates.

a. Identify the root aggregates: `storage aggregate show -root true`

b. Display the pool1 aggregates: `storage aggregate plex show -pool 1`

c. Offline and delete the remote plex of the root aggregate:

```
aggr plex offline <root-aggregate> -plex <remote-plex-for-root-aggregate>
```

```
aggr plex delete <root-aggregate> -plex <remote-plex-for-root-aggregate>
```

For example:

```
# aggr plex offline aggr0_node_A_1-FC_01 -plex remoteplex4
```

```
# aggr plex delete aggr0_node_A_1-FC_01 -plex remoteplex4
```

4. Confirm the mailbox count, disk autoassign, and transition mode before proceeding using the following commands on each controller:

a. Set the advanced privilege level: `set -priv advanced`

b. Confirm that only three mailbox drives are shown for each controller module: `storage failover mailbox-disk show`

c. Return to admin mode: `set -priv admin`

d. Confirm that the transition mode is disruptive: `metrocluster transition show`

5. Check for any broken disks: `disk show -broken`

6. Remove or replace any broken disks

7. Confirm aggregates are healthy by using the following commands on node_A_1-FC and node_B_1-FC:

```
storage aggregate show
```

```
storage aggregate plex show
```

The `storage aggregate show` command indicates that the root aggregate is unmirrored.

8. Check for any VLANs or interface groups:

```
network port ifgrp show
```

```
network port vlan show
```

If none are present, skip the following two steps.

9. Display the list of Lifs using VLANs or ifgrps:

```
network interface show -fields home-port,curr-port
```

```
network port show -type if-group | vlan
```

10. Remove any VLANs and interface groups.

You must perform these steps for all LIFs in all SVMs, including those SVMs with the -mc suffix.

- a. Move any LIFs using the VLANs or interface groups to an available port: `network interface modify -vserver vservice-name -lif lif_name -home- port port`
- b. Display the LIFs that are not on their home ports: `network interface show -is-home false`
- c. Revert all LIFs to their respective home ports: `network interface revert -vserver vservice_name -lif lif_name`
- d. Verify that all LIFs are on their home ports: `network interface show -is-home false`

No LIFs should appear in the output.

- e. Remove VLAN and ifgrp ports from broadcast domain: `network port broadcast-domain remove-ports -ip-space ip-space -broadcast-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..`
- f. Verify that all the vlan and ifgrp ports are not assigned to a broadcast domain: `network port show -type if-group | vlan`
- g. Delete all VLANs: `network port vlan delete -node nodename -vlan-name vlan-name`
- h. Delete interface groups: `network port ifgrp delete -node nodename -ifgrp ifgrp-name`

11. Move any LIFs as required to resolve conflicts with the MetroCluster IP interface ports.

You must move the LIFs identified in step 1 of [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#).

- a. Move any LIFs hosted on the desired port to another port: `network interface modify -lif lifname -vserver vservice-name -home-port new-homeport`network interface revert -lif lifname -vserver vservice-name`
- b. If necessary, move the destination port to an appropriate IPspace and broadcast domain. `network port broadcast-domain remove-ports -ip-space current-ip-space -broadcast-domain current-broadcast-domain -ports controller-name:current-port`network port broadcast-domain add-ports -ip-space new-ip-space -broadcast-domain new-broadcast-domain -ports controller-name:new-port`

12. Halt the MetroCluster FC controllers (node_A_1-FC and node_B_1-FC): `system node halt`

13. At the LOADER prompt, synchronize the hardware clocks between the FC and IP controller modules.

- a. On the old MetroCluster FC node (node_A_1-FC), display the date: `show date`
- b. On the new MetroCluster IP controllers (node_A_1-IP and node_B_1-IP), set the date shown on original controller: `set date mm/dd/yy`
- c. On the new MetroCluster IP controllers (node_A_1-IP and node_B_1-IP), verify the date: `show date`

14. Halt and power off the MetroCluster FC controller modules (node_A_1-FC and node_B_1-FC), FC-to-SAS bridges (if present), FC switches (if present) and each storage shelf connected to these nodes.

15. Disconnect the shelves from the MetroCluster FC controllers and document which shelves are local storage to each cluster.
16. If the configuration uses FC-to-SAS bridges or FC back-end switches, disconnect and remove them.

Remove FC-to-SAS bridges

a. Identify the bridges:

```
system bridge show
```

b. Remove the bridges:

```
system bridge remove -name <bridge_name>
```

c. Confirm the bridges are removed:

```
system bridge show
```

The following example shows that the bridges are removed:

Example

```
cluster1::> system bridge remove -name ATTO_10.226.197.16
cluster1::> system bridge show

Is          Monitor
  Bridge    Symbolic Name Vendor  Model      Bridge WWN
Monitored Status
-----
-----
      ATTO_FibreBridge6500N_1
                Bridge Number 16
                        Atto    FibreBridge 6500N
                                2000001086603824
false      -
      ATTO_FibreBridge6500N_2
                Not Set      Atto    FibreBridge 6500N
                                20000010866037e8
false      -
      ATTO_FibreBridge6500N_3
                Not Set      Atto    FibreBridge 6500N
                                2000001086609e0e
false      -
      ATTO_FibreBridge6500N_4
                Not Set      Atto    FibreBridge 6500N
                                2000001086609c06
false      -
      4 entries were displayed.
```

Remove FC switches

a. Identify the switches:

```
system switch fibre-channel show
```

b. Remove the switches:

```
system switch fibre-channel remove -switch-name <switch_name>
```

c. Confirm the switches are removed:

```
system switch fibre-channel show
```

Example

```
cluster1::> system switch fibre-channel show
          Symbolic                               Is
Monitor
  Switch      Name      Vendor  Model      Switch WWN
Monitored Status
-----
Cisco_10.226.197.34
          mcc-cisco-8Gb-fab-4
                  Cisco  DS-C9148-16P-K9
                              2000547fee78f088
true      ok
          mcc-cisco-8Gb-fab-1
                  mcc-cisco-8Gb-fab-1
                          Cisco  -
false     -
          mcc-cisco-8Gb-fab-2
                  mcc-cisco-8Gb-fab-2
                          Cisco  -
false     -
          mcc-cisco-8Gb-fab-3
                  mcc-cisco-8Gb-fab-3
                          Cisco  -
false     -
          4 entries were displayed.
cluster1::> system switch fibre-channel remove -switch-name
Cisco_10.226.197.34
cluster1::> system switch fibre-channel show
          Symbolic                               Is
Monitor
  Switch      Name      Vendor  Model      Switch WWN
Monitored Status
-----
          mcc-cisco-8Gb-fab-4
                  mcc-cisco-8Gb-fab-4
                          Cisco
                              -
false     -
          mcc-cisco-8Gb-fab-1
                  mcc-cisco-8Gb-fab-1
                          Cisco  -
false     -
          mcc-cisco-8Gb-fab-2
```

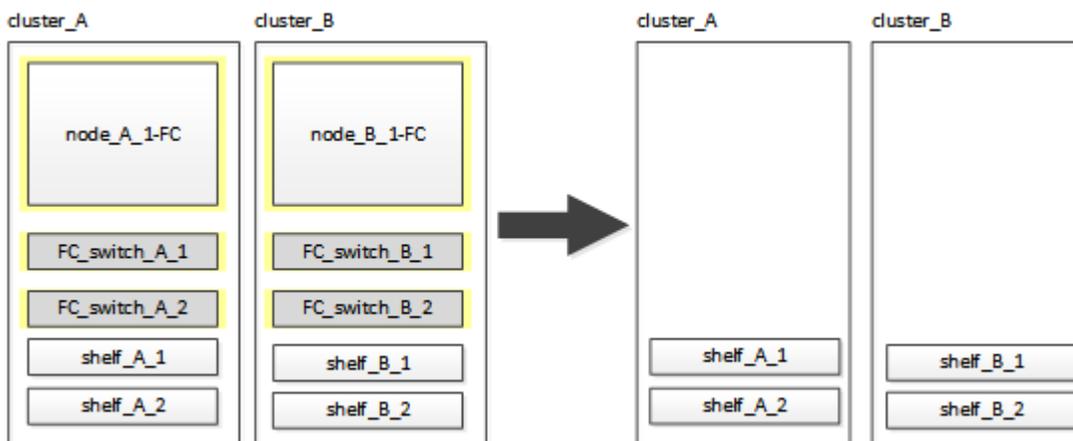
```

      mcc-cisco-8Gb-fab-2
          Cisco - -
false -
      mcc-cisco-8Gb-fab-3
          mcc-cisco-8Gb-fab-3
          Cisco - -
false -
      4 entries were displayed
cluster1::>

```

17. In Maintenance mode on the MetroCluster FC nodes (node_A_1-FC and node_B_1-FC), confirm no disks are connected: `disk show -v`
18. Power down and remove the MetroCluster FC nodes.

At this point, the MetroCluster FC controllers have been removed and the shelves are disconnected from all controllers.



Connecting the MetroCluster IP controller modules

You must add the four new controller modules and any additional storage shelves to the configuration. The new controller modules are added two-at-a-time.

Setting up the new controllers

You must rack and cable the new MetroCluster IP controllers to the storage shelves previously connected to the MetroCluster FC controllers.

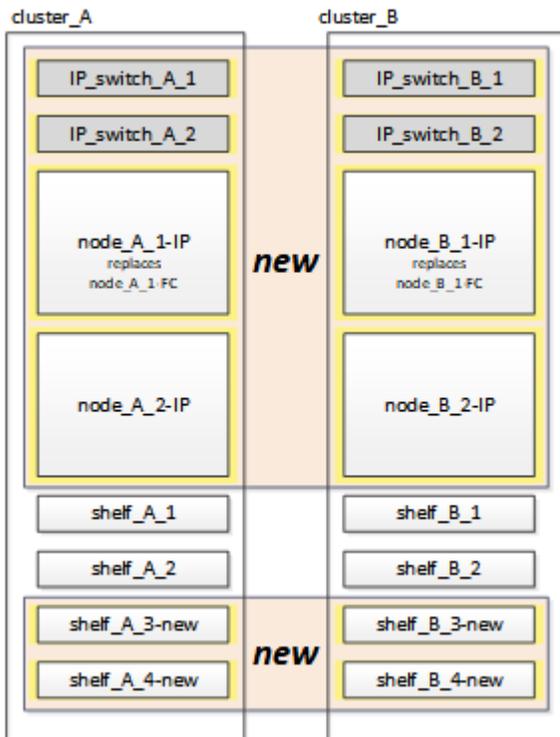
About this task

These steps must be performed on each of the MetroCluster IP nodes.

- node_A_1-IP
- node_A_2-IP
- node_B_1-IP

- node_B_2-IP

In the following example, two additional storage shelves are added at each site to provide storage to accommodate the new controller modules.



Steps

1. Plan out the positioning of the new controller modules and storage shelves as needed.

The rack space depends on the platform model of the controller modules, the switch types, and the number of storage shelves in your configuration.

2. Properly ground yourself.
3. Rack the new equipment: controllers, storage shelves, and IP switches.

Do not cable the storage shelves or IP switches at this time.

4. Connect the power cables and management console connection to the controllers.
5. Verify that all storage shelves are powered off.
6. Verify that no drives are connected by performing the following steps on all four nodes:

- a. At the LOADER prompt, launch the boot menu:

```
boot_ontap maint
```

- b. Verify that no drives are connected:

```
disk show -v
```

The output should show no drives.

- c. Halt the node:

```
halt
```

7. Boot all four nodes using the 9a option on the boot menu.

a. At the LOADER prompt, launch the boot menu:

```
boot_ontap menu
```

b. At the boot menu, select option “9a” to reboot the controller.

c. Let the controller module complete booting before moving to the next controller module.

After “9a” completes, the nodes automatically return to the boot menu.

8. Cable the storage shelves.

Refer to the controller installation and setup procedures for your model for cabling information.

[ONTAP Hardware Systems Documentation](#)

9. Cable the controllers to the IP switches as described in [Cabling the IP switches](#).

10. Prepare the IP switches for the application of the new RCF files.

Follow the steps for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

11. Download and install the RCF files.

Follow the steps for your switch vendor:

- [Download and install the Broadcom IP RCF files](#)
- [Download and install the Cisco IP RCF files](#)
- [Download and install the NVIDIA RCF files\]](#)

12. Turn on power to the first new controller (node_A_1-IP) and press Ctrl-C to interrupt the boot process and display the LOADER prompt.

13. Boot the controller to Maintenance mode:

```
boot_ontap_maint
```

14. Display the system ID for the controller:

```
sysconfig -v
```

15. Confirm that the shelves from the existing configuration are visible from the new MetroCluster IP node:

```
storage show shelf``disk show -v
```

16. Halt the node:

```
halt
```

17. Repeat the preceding steps on the other node at the partner site (site_B).

Connecting and booting up node_A_1-IP and node_B_1-IP

After connecting the MetroCluster IP controllers and IP switches, you transition and boot up node_A_1-IP and node_B_1-IP.

Bringing up node_A_1-IP

You must boot the node with the correct transition option.

Steps

1. Boot node_A_1-IP to the boot menu:

```
boot_ontap menu
```

2. Issue the following command at the boot menu prompt to initiate transition:

```
boot_after_mcc_transition
```

- This command reassigns all the disks owned by node_A_1-FC to node_A_1-IP.
 - node_A_1-FC disks are assigned to node_A_1-IP
 - node_B_1-FC disks are assigned to node_B_1-IP
- The command also automatically makes other required system ID reassignments so the MetroCluster IP nodes can boot to the ONTAP prompt.
- If the boot_after_mcc_transition command fails for any reason, it should be re-run from the boot menu.



- If the following prompt is displayed, enter Ctrl-C to continue. Checking MCC DR state... [enter Ctrl-C(resume), S(status), L(link)]_
- If the root volume was encrypted, the node halts with the following message. Halting the system, because root volume is encrypted (NetApp Volume Encryption) and the key import failed. If this cluster is configured with external (KMIP) key-manager, check the health of the key servers.

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning. Selection (1-9)?

```
`boot_after_mcc_transition`
```

```
This will replace all flash-based configuration with the last backup  
to disks. Are you sure you want to continue?: yes
```

```
MetroCluster Transition: Name of the MetroCluster FC node: `node_A_1-  
FC`
```

```
MetroCluster Transition: Please confirm if this is the correct value  
[yes|no]:? y
```

```
MetroCluster Transition: Disaster Recovery partner sysid of  
MetroCluster FC node node_A_1-FC: `systemID-of-node_B_1-FC`
```

```
MetroCluster Transition: Please confirm if this is the correct value  
[yes|no]:? y
```

```
MetroCluster Transition: Disaster Recovery partner sysid of local  
MetroCluster IP node: `systemID-of-node_B_1-IP`
```

```
MetroCluster Transition: Please confirm if this is the correct value  
[yes|no]:? y
```

3. If data volumes are encrypted, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> <p>For more information, see Restoring onboard key management encryption keys.</p>
External key management	<pre>security key-manager key query -node node-name</pre> <p>For more information, see Restoring external key management encryption keys.</p>

4. If the root volume is encrypted, use the procedure in [Recovering key management if the root volume is encrypted](#).

Recovering key management if the root volume is encrypted

If the root volume is encrypted, you must use special boot commands to restore the key management.

Before you begin

You must have the passphrases gathered earlier.

Steps

1. If onboard key management is used, perform the following substeps to restore the configuration.
 - a. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- b. Select option "(10) Set onboard key management recovery secrets" from the boot menu.

Respond as appropriate to the prompts:

```
This option must be used only in disaster recovery procedures. Are
you sure? (y or n): y
Enter the passphrase for onboard key management: passphrase
Enter the passphrase again to confirm: passphrase

Enter the backup data: backup-key
```

The system boots to the boot menu.

- c. Enter option "6" at the boot menu.

Respond as appropriate to the prompts:

```
This will replace all flash-based configuration with the last backup
to
disks. Are you sure you want to continue?: y

Following this, the system will reboot a few times and the following
prompt will be available continue by saying y

WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
```

After the reboots, the system will be at the LOADER prompt.

- d. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- e. Again elect option "(10) Set onboard key management recovery secrets" from the boot menu.

Respond as appropriate to the prompts:

```
This option must be used only in disaster recovery procedures. Are
you sure? (y or n): `y`
Enter the passphrase for onboard key management: `passphrase`
Enter the passphrase again to confirm: `passphrase`

Enter the backup data: `backup-key`
```

The system boots to the boot menu.

f. Enter option “1” at the boot menu.

If the following prompt is displayed, you can press Ctrl+C to resume the process.

```
Checking MCC DR state... [enter Ctrl-C(resume), S(status), L(link)]
```

The system boots to the ONTAP prompt.

g. Restore the onboard key management:

```
security key-manager onboard sync
```

Respond as appropriate to the prompts, using the passphrase you collected earlier:

```
cluster_A::> security key-manager onboard sync
Enter the cluster-wide passphrase for onboard key management in
Vserver "cluster_A":: passphrase
```

2. If external key management is used, perform the following substeps to restore the configuration.

a. Set the required bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address

setenv bootarg.kmip.init.netmask netmask

setenv bootarg.kmip.init.gateway gateway-address

setenv bootarg.kmip.init.interface interface-id
```

b. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

c. Select option “(11) Configure node for external key management” from the boot menu.

The system boots to the boot menu.

- d. Enter option “6” at the boot menu.

The system boots multiple times. You can respond affirmatively when prompted to continue the boot process.

After the reboots, the system will be at the LOADER prompt.

- e. Set the required bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address  
  
setenv bootarg.kmip.init.netmask netmask  
  
setenv bootarg.kmip.init.gateway gateway-address  
  
setenv bootarg.kmip.init.interface interface-id
```

- f. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- g. Again select option “(11) Configure node for external key management” from the boot menu and respond to the prompts as required.

The system boots to the boot menu.

- h. Restore the external key management:

```
security key-manager external restore
```

Creating the network configuration

You must create a network configuration that matches the configuration on the FC nodes. This is because the MetroCluster IP node replays the same configuration when it boots, which means that when node_A_1-IP and node_B_1-IP boot, ONTAP will try to host LIFs on the same ports that were used on node_A_1-FC and node_B_1-FC respectively.

About this task

As you create the network configuration, use the plan made in [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#) to assist you.



Additional configuration may be needed to bring up data LIFs after the MetroCluster IP nodes have been configured.

Steps

1. Verify that all cluster ports are in the appropriate broadcast domain:

The cluster IPspace and cluster broadcast domain are required in order to create cluster LIFs

- a. View the IP spaces:

```
network ipspace show
```

- b. Create IP spaces and assign cluster ports as needed.

[Configuring IPspaces \(cluster administrators only\)](#)

- c. View the broadcast domains:

```
network port broadcast-domain show
```

- d. Add any cluster ports to a broadcast domain as needed.

[Adding or removing ports from a broadcast domain](#)

- e. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Creating a VLAN](#)

[Combining physical ports to create interface groups](#)

2. Verify that MTU settings are set correctly for the ports and broadcast domain and make changes using the following commands:

```
network port broadcast-domain show
```

```
network port broadcast-domain modify -broadcast-domain bcastdomainname -mtu mtu-value
```

Setting up cluster ports and cluster LIFs

You must set up cluster ports and LIFs. The following steps need to be performed on the site A nodes which were booted up with root aggregates.

Steps

1. Identify the list of LIFs using the desired Cluster port:

```
network interface show -curr-port portname
```

```
network interface show -home-port portname
```

2. For each cluster port, change the home port of any of the LIFs on that port to another port,
 - a. Enter advanced privilege mode and enter “y” when prompted to continue:

```
set priv advanced
```

- b. If the LIF being modified is a data LIF:

```
vserver config override -command "network interface modify -lif lifname -vserver vservername -home-port new-datahomeport"
```

- c. If the LIF is not a data LIF:

```
network interface modify -lif lifname -vserver vservername -home-port new-
```

datahomeport

- d. Revert the modified LIFs to their home port:

```
network interface revert * -vserver vserver_name
```

- e. Verify that there are no LIFs on the cluster port:

```
network interface show -curr-port portname
```

```
network interface show -home-port portname
```

- f. Remove the port from the current broadcast domain:

```
network port broadcast-domain remove-ports -ipSPACE ipSPACEname -broadcast-domain bcastdomainname -ports node_name:port_name
```

- g. Add the port to the cluster IPspace and broadcast domain:

```
network port broadcast-domain add-ports -ipSPACE Cluster -broadcast-domain Cluster -ports node_name:port_name
```

- h. Verify that the port's role has changed: `network port show`

- i. Repeat these substeps for each cluster port.

- j. Return to admin mode:

```
set priv admin
```

3. Create cluster LIFs on the new cluster ports:

- a. For autoconfiguration using link-local address for cluster LIF, use the following command:

```
network interface create -vserver Cluster -lif cluster_lifname -service -policy default-cluster -home-node a1name -home-port clusterport -auto true
```

- b. To assign static IP address for the cluster LIF, use the following command:

```
network interface create -vserver Cluster -lif cluster_lifname -service -policy default-cluster -home-node a1name -home-port clusterport -address ip-address -netmask netmask -status-admin up
```

Verifying LIF configuration

The node management LIF, cluster management LIF and intercluster LIF will still be present after the storage movement from the old controller. If necessary, you must move LIFs to appropriate ports.

Steps

1. Verify whether the management LIF and cluster management LIFs are on the desired port already:

```
network interface show -service-policy default-management
```

```
network interface show -service-policy default-intercluster
```

If the LIFs are on the desired ports, you can skip the rest of the steps in this task and proceed to the next task.

2. For each node, cluster management, or intercluster LIFs that are not on the desired port, change the home port of any of the LIFs on that port to another port.

- a. Repurpose the desired port by moving any LIFs hosted on desired port to another port:

```
vserver config override -command "network interface modify -lif lifname
-vserver vservername -home-port new-datahomeport"
```

- b. Revert the modified LIFs to their new home port:

```
vserver config override -command "network interface revert -lif lifname
-vserver _vservername"
```

- c. If the desired port is not in the right IPspace and broadcast domain, remove the port from the current IPspace and broadcast domain:

```
network port broadcast-domain remove-ports -ip-space current-ip-space
-broadcast-domain current-broadcast-domain -ports controller-name:current-
port
```

- d. Move the desired port to the right IPspace and broadcast domain:

```
network port broadcast-domain add-ports -ip-space new-ip-space -broadcast
-domain new-broadcast-domain -ports controller-name:new-port
```

- e. Verify that the port's role has changed:

```
network port show
```

- f. Repeat these substeps for each port.

3. Move node, cluster management LIFs, and intercluster LIF to the desired port:

- a. Change the LIF's home port:

```
network interface modify -vserver vserver -lif node_mgmt -home-port port
-home-node homenode
```

- b. Revert the LIF to its new home port:

```
network interface revert -lif node_mgmt -vserver vservername
```

- c. Change the cluster management LIF's home port:

```
network interface modify -vserver vserver -lif cluster-mgmt-LIF-name -home
-port port -home-node homenode
```

- d. Revert the cluster management LIF to its new home port:

```
network interface revert -lif cluster-mgmt-LIF-name -vserver vservername
```

- e. Change the intercluster LIF's home port:

```
network interface modify -vserver vsverver -lif intercluster-lif-name -home
-node nodename -home-port port
```

f. Revert the intercluster LIF to its new home port:

```
network interface revert -lif intercluster-lif-name -vserver vsververname
```

Bringing up node_A_2-IP and node_B_2-IP

You must bring up and configure the new MetroCluster IP node at each site, creating an HA pair in each site.

Bringing up node_A_2-IP and node_B_2-IP

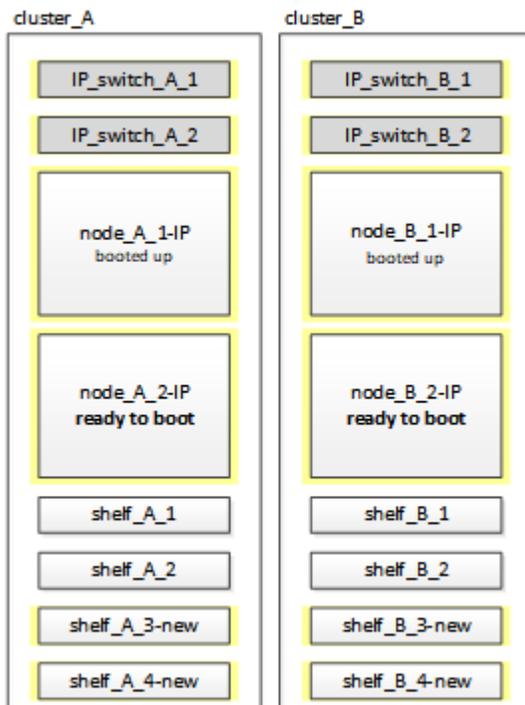
You must boot the new controller modules one at a time using the correct option at the boot menu.

About this task

In these steps, you boot up the two brand new nodes, expanding what had been a two-node configuration into a four-node configuration.

These steps are performed on the following nodes:

- node_A_2-IP
- node_B_2-IP



Steps

1. Boot the new nodes using boot option "9c".

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning. Selection (1-9)? 9c

The node initializes and boots to the node setup wizard, similar to the following.

```
Welcome to node setup
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
Any changes you made before quitting will be saved.
To accept a default or omit a question, do not enter a value. .
.
.
```

If option "9c" does not succeed, take the following steps to avoid possible data loss:

- Do not attempt to run option 9a.
- Physically disconnect the existing shelves that contain data from the original MetroCluster FC configuration (shelf_A_1, shelf_A_2, shelf_B_1, shelf_B_2).
- Contact technical support, referencing the KB article [MetroCluster FC to IP transition - Option 9c Failing](#).

[NetApp Support](#)

2. Enable the AutoSupport tool by following the directions provided by the wizard.
3. Respond to the prompts to configure the node management interface.

```
Enter the node management interface port: [e0M]:
Enter the node management interface IP address: 10.228.160.229
Enter the node management interface netmask: 225.225.252.0
Enter the node management interface default gateway: 10.228.160.1
```

4. Verify that the storage failover mode is set to HA:

```
storage failover show -fields mode
```

If the mode is not HA, set it:

```
storage failover modify -mode ha -node localhost
```

You must then reboot the node for the change to take effect.

5. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```
cluster01::> network port show
```

(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed Admin/Oper

cluster01-01						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
cluster01-02						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000

6. Exit the Node Setup wizard:

```
exit
```

7. Log into the admin account using the admin user name.

8. Join the existing cluster using the Cluster Setup wizard.

```

:> cluster setup
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and "exit"
or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{create, join}:
join

```

9. After you complete the Cluster Setup wizard and it exits, verify that the cluster is active and the node is healthy:

```
cluster show
```

10. Disable disk autoassignment:

```
storage disk option modify -autoassign off -node node_A_2-IP
```

11. If encryption is used, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> <p>For more information, see Restoring onboard key management encryption keys.</p>
External key management	<pre>security key-manager key query -node node-name</pre> <p>For more information, see Restoring external key management encryption keys.</p>

12. Repeat the above steps on the second new controller module (node_B_2-IP).

Verifying MTU settings

Verify that MTU settings are set correctly for the ports and broadcast domain and make changes.

Steps

1. Check the MTU size used in the cluster broadcast domain:

```
network port broadcast-domain show
```

2. If necessary, update the MTU size as needed:

```
network port broadcast-domain modify -broadcast-domain bcast-domain-name -mtu
mtu-size
```

Configuring intercluster LIFs

Configure the intercluster LIFs required for cluster peering.

This task must be performed on both of the new nodes, `node_A_2-IP` and `node_B_2-IP`.

Step

1. Configure the intercluster LIFs. See [Configuring intercluster LIFs](#)

Verifying cluster peering

Verify that `cluster_A` and `cluster_B` are peered and nodes on each cluster can communicate with each other.

Steps

1. Verify the cluster peering relationship:

```
cluster peer health show
```

```
cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
node_A_1-IP
          cluster_B          node_B_1-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          node_B_2-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
node_A_2-IP
image:../media/transition_2n_booting_a_2_and_b_2.png["Booting new IP
nodes during transition"]
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          node_B_2-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
```

2. Ping to check that the peer addresses are reachable:

```
cluster peer ping -originating-node local-node -destination-cluster remote-
cluster-name
```

Configuring the new nodes and completing transition

With the new nodes added, you must complete the transition steps and configure the MetroCluster IP nodes.

Configuring the MetroCluster IP nodes and disabling transition

You must implement the MetroCluster IP connections, refresh the MetroCluster configuration, and disable transition mode.

Steps

1. Form the new nodes into a DR group by issuing the following commands from controller node_A_1-IP:

```
metrocluster configuration-settings dr-group create -partner-cluster  
<peer_cluster_name> -local-node <local_controller_name> -remote-node  
<remote_controller_name>
```

```
metrocluster configuration-settings dr-group show
```

2. Create MetroCluster IP interfaces (node_A_1-IP, node_A_2-IP, node_B_1-IP, node_B_2-IP) — two interfaces need to be created per controller; eight interfaces in total:



Do not use 169.254.17.x or 169.254.18.x IP addresses when you create MetroCluster IP interfaces to avoid conflicts with system auto-generated interface IP addresses in the same range.

```
metrocluster configuration-settings interface create -cluster-name  
<cluster_name> -home-node <controller_name> -home-port <port_name> -address  
<ip_address> -netmask <netmask_address> -vlan-id <vlan_id>
```

```
metrocluster configuration-settings interface show
```

Certain platforms use a VLAN for the MetroCluster IP interface. By default, each of the two ports use a different VLAN: 10 and 20.

If supported, you can also specify a different (non-default) VLAN higher than 100 (between 101 and 4095) using the `-vlan-id` parameter in the `metrocluster configuration-settings interface create` command.

The following platforms do **not** support the `-vlan-id` parameter:

- FAS8200 and AFF A300
- AFF A320
- FAS9000 and AFF A700
- AFF C800, ASA C800, AFF A800 and ASA A800

All other platforms support the `-vlan-id` parameter.

The default and valid VLAN assignments depend on whether the platform supports the `-vlan-id` parameter:

Platforms that support `-vlan-id`

Default VLAN:

- When the `-vlan-id` parameter is not specified, the interfaces are created with VLAN 10 for the "A" ports and VLAN 20 for the "B" ports.
- The VLAN specified must match the VLAN selected in the RCF.

Valid VLAN ranges:

- Default VLAN 10 and 20
- VLANs 101 and higher (between 101 and 4095)

Platforms that do not support `-vlan-id`

Default VLAN:

- Not applicable. The interface does not require a VLAN to be specified on the MetroCluster interface. The switch port defines the VLAN that is used.

Valid VLAN ranges:

- All VLANs not explicitly excluded when generating the RCF. The RCF alerts you if the VLAN is invalid.

3. Perform the MetroCluster connect operation from controller node_A_1-IP to connect the MetroCluster sites — this operation can take a few minutes to complete:

```
metrocluster configuration-settings connection connect
```

4. Verify that the remote cluster disks are visible from each controller via the iSCSI connections:

```
disk show
```

You should see the remote disks belonging to the other nodes in the configuration.

5. Mirror the root aggregate for node_A_1-IP and node_B_1-IP:

```
aggregate mirror -aggregate root-aggr
```

6. Assign disks for node_A_2-IP and node_B_2-IP.

Pool 1 disk assignments were already made for node_A_1-IP and node_B_1-IP when the `boot_after_mcc_transition` command was issued at the boot menu.

- a. Issue the following commands on node_A_2-IP:

```
disk assign disk1disk2disk3 ... diskn -sysid node_B_2-IP-controller-sysid  
-pool 1 -force
```

- b. Issue the following commands on node_B_2-IP:

```
disk assign disk1disk2disk3 ... diskn -sysid node_A_2-IP-controller-sysid
```

```
-pool 1 -force
```

7. Confirm ownership has been updated for the remote disks:

```
disk show
```

8. If necessary, refresh the ownership information using the following commands:

a. Go to advanced privilege mode and enter y when prompted to continue:

```
set priv advanced
```

b. Refresh disk ownership:

```
disk refresh-ownership controller-name
```

c. Return to admin mode:

```
set priv admin
```

9. Mirror the root aggregates for node_A_2-IP and node_B_2-IP:

```
aggregate mirror -aggregate root-aggr
```

10. Verify that the aggregate re-synchronization has completed for root and data aggregates:

```
aggr show` `aggr plex show
```

The resync can take some time but must complete before proceeding with the following steps.

11. Refresh the MetroCluster configuration to incorporate the new nodes:

a. Go to advanced privilege mode and enter y when prompted to continue:

```
set priv advanced
```

b. Refresh the configuration:

If you have configured...	Issue this command...
A single aggregate in each cluster:	<pre>metrocluster configure -refresh true -allow-with-one-aggregate true</pre>
More than a single aggregate in each cluster	<pre>metrocluster configure -refresh true</pre>

c. Return to admin mode:

```
set priv admin
```

12. Disable MetroCluster transition mode:

a. Enter advanced privilege mode and enter “y” when prompted to continue:

```
set priv advanced
```

- b. Disable transition mode:

```
metrocluster transition disable
```

- c. Return to admin mode:

```
set priv admin
```

Setting up data LIFs on the new nodes

You must configure data LIFs on the new nodes, node_A_2-IP and node_B_2-IP.

You must add any new ports available on new controllers to a broadcast domain if not already assigned to one. If required, create VLANs or interface groups on the new ports. See [Network management](#)

1. Identify the current port usage and broadcast domains:

```
network port show ``network port broadcast-domain show
```

2. Add ports to broadcast domains and VLANs as necessary.

- a. View the IP spaces:

```
network ipspace show
```

- b. Create IP spaces and assign data ports as needed.

[Configuring IPspaces \(cluster administrators only\)](#)

- c. View the broadcast domains:

```
network port broadcast-domain show
```

- d. Add any data ports to a broadcast domain as needed.

[Adding or removing ports from a broadcast domain](#)

- e. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Creating a VLAN](#)

[Combining physical ports to create interface groups](#)

3. Verify that the LIFs are hosted on the appropriate node and ports on the MetroCluster IP nodes (including the SVM with -mc vserver) as needed.

See the information gathered in [Creating the network configuration](#).

- a. Check the home port of the LIFs:

```
network interface show -field home-port
```

- b. If necessary, modify the LIF configuration:

```
vserver config override -command "network interface modify -vserver
<svm_name> -home-port <active_port_after_upgrade> -lif <lif_name> -home-node
<new_node_name>
```

- c. Revert the LIFs to their home ports:

```
network interface revert * -vserver <svm_name>
```

Bringing up the SVMs

Due to the changes if LIF configuration, you must restart the SVMs on the new nodes.

Steps

1. Check the state of the SVMs:

```
metrocluster vserver show
```

2. Restart the SVMs on cluster_A that do not have an “-mc” suffix:

```
vserver start -vserver <svm_name> -force true
```

3. Repeat the previous steps on the partner cluster.
4. Check that all SVMs are in a healthy state:

```
metrocluster vserver show
```

5. Verify that all data LIFs are online:

```
network interface show
```

Moving a system volume to the new nodes

To improve resiliency, a system volume should be moved from controller node_A_1-IP to controller node_A_2-IP, and also from node_B_1-IP to node_B_2-IP. You must create a mirrored aggregate on the destination node for the system volume.

About this task

System volumes have the name form “MDV_CRS_*_A” or “MDV_CRS_*_B.” The designations “_A” and “_B” are unrelated to the site_A and site_B references used throughout this section; e.g., MDV_CRS_*_A is not associated with site_A.

Steps

1. Assign at least three pool 0 and three pool 1 disks each for controllers node_A_2-IP and node_B_2-IP as needed.
2. Enable disk auto-assignment.
3. Move the _B system volume from node_A_1-IP to node_A_2-IP using the following steps from site_A.
 - a. Create a mirrored aggregate on controller node_A_2-IP to hold the system volume:

```
aggr create -aggregate new_node_A_2-IP_aggr -diskcount 10 -mirror true -node
node_A_2-IP
```

```
aggr show
```

The mirrored aggregate requires five pool 0 and five pool 1 spare disks owned by controller node_A_2-IP.

The advanced option, “-force-small-aggregate true” can be used to limit disk use to 3 pool 0 and 3 pool 1 disks, if disks are in short supply.

- b. List the system volumes associated with the admin SVM:

```
vserver show
```

```
volume show -vserver <admin_svm_name>
```

You should identify volumes contained by aggregates owned by site_A. The site_B system volumes will also be shown.

4. Move the MDV_CRS_*_B system volume for site_A to the mirrored aggregate created on controller node_A_2-IP

- a. Check for possible destination aggregates:

```
volume move target-aggr show -vserver <admin_svm_name> -volume MDV_CRS_*_B
```

The newly created aggregate on node_A_2-IP should be listed.

- b. Move the volume to the newly created aggregate on node_A_2-IP:

```
set advanced
```

```
volume move start -vserver <admin_svm_name> -volume MDV_CRS_*_B -destination  
-aggregate new_node_A_2-IP_aggr -cutover-window 40
```

- c. Check status for the move operation:

```
volume move show -vserver <admin_svm_name> -volume MDV_CRS_*_B
```

- d. When the move operation complete, verify that the MDV_CRS_*_B system is contained by the new aggregate on node_A_2-IP:

```
set admin
```

```
volume show -vserver <admin_svm_name>
```

5. Repeat the above steps on site_B (node_B_1-IP and node_B_2-IP).

Returning the system to normal operation

You must perform final configuration steps and return the MetroCluster configuration to normal operation.

Verifying MetroCluster operation and assigning drives after transition

You must verify that the MetroCluster is operating correctly and assign drives to the second pair of new nodes

(node_A_2-IP and node_B_2-IP).

1. Confirm that the MetroCluster configuration-type is IP-fabric: `metrocluster show`
2. Perform a MetroCluster check.
 - a. Issue the following command: `metrocluster check run`
 - b. Display the results of the MetroCluster check: `metrocluster check show`
3. Confirm that the DR group with the MetroCluster IP nodes is configured: `metrocluster node show`
4. Create and mirror additional data aggregates for controllers node_A_2-IP and node_B_2-IP at each site as needed.

Installing licenses for the new controller module

You must add licenses for the new controller module for any ONTAP services that require standard (node-locked) licenses. For features with standard licenses, each node in the cluster must have its own key for the feature.

For detailed information about licensing, see the knowledgebase article 3013749: Data ONTAP 8.2 Licensing Overview and References on the NetApp Support Site and the *System Administration Reference*.

1. If necessary, obtain license keys for the new node on the NetApp Support Site in the My Support section under Software licenses.

For further information on license replacements, see the Knowledge Base article [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#).

2. Issue the following command to install each license key: `system license add -license-code license_key`

The `license_key` is 28 digits in length.

Repeat this step for each required standard (node-locked) license.

Completing configuration of the nodes

There are miscellaneous configuration steps that can be performed prior to completing the procedures. Some of these steps are optional.

1. Configure the service processor: `system service-processor network modify`
2. Set up autosupport on the new nodes: `system node autosupport modify`
3. The controllers can be optionally renamed as part of the transition. The following command is used to rename a controller: `system node rename -node <old-name> -newname <new-name>`

The renaming operation can take a few minutes to complete. Confirm that any name changes have propagated to each node prior to continuing with other steps using the `system show -fields node` command.

4. Configure a monitoring service as desired.

[Considerations for Mediator](#)

Sending a custom AutoSupport message after maintenance

After completing the transition, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

1. To resume automatic support case generation, send an AutoSupport message to indicate that the maintenance is complete.
 - a. Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=end`
 - b. Repeat the command on the partner cluster.

Disruptively transitioning from MetroCluster FC to MetroCluster IP when retiring storage shelves (ONTAP 9.8 and later)

Beginning with ONTAP 9.8, you can disruptively transition a two-node MetroCluster FC configuration to a four-node MetroCluster IP configuration and retire the existing storage shelves. The procedure includes steps to move data from the existing drive shelves to the new configuration, and then retire the old shelves.

- This procedure is used when you plan to retire the existing storage shelves and move all data to the new shelves in the MetroCluster IP configuration.
- The existing storage shelf models must be supported by the new MetroCluster IP nodes.
- This procedure is supported on systems running ONTAP 9.8 and later.
- This procedure is disruptive.
- This procedure applies only to a two-node MetroCluster FC configuration.

If you have a four-node MetroCluster FC configuration, see [Choosing your transition procedure](#).

- You must meet all requirements and follow all steps in the procedure.

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the

"Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Requirements for transition when retiring old shelves

Before starting the transition process, you must make sure the existing MetroCluster FC configuration meets the requirements.

- It must be a two-node fabric-attached or stretch MetroCluster configuration and all nodes must be running ONTAP 9.8 or later.

The new MetroCluster IP controller modules should be running the same version of ONTAP 9.8.

- The existing and new platforms must be a supported combination for transition.

[Supported platforms for nondisruptive transition](#)

- It must meet all requirements and cabling as described in the *MetroCluster Installation and Configuration Guides*.

[Fabric-attached MetroCluster installation and configuration](#)

The new configuration must also meet the following requirements:

- The new MetroCluster IP platform models must support the old storage shelf models.

[NetApp Hardware Universe](#)

- Depending on the spare disks available in the existing shelves, additional drives must be added.

This might require additional drive shelves.

You need to have additional 14 to 18 drives for each controller:

- Three pool 0 drives
 - Three pool 1 drives
 - Two spare drives
 - Six to ten drives for the system volume
- You must ensure that the configuration, including the new nodes, does not exceed the platform limits for the configuration, including drive count, root aggregate size capacity, etc.

This information is available for each platform model at [NetApp Hardware Universe](#)

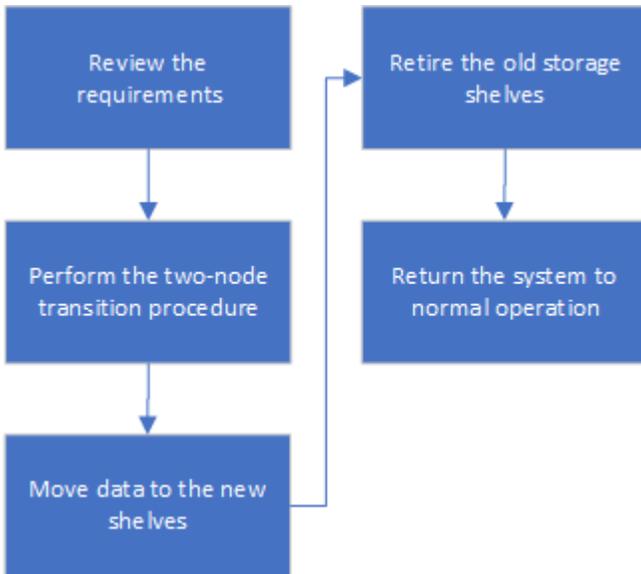
You must have remote console access for all six nodes from either MetroCluster site or plan for travel between the sites as required by the procedure.

Workflow for disruptive transition when moving data and retiring old storage shelves

You must follow the specific workflow to ensure a successful transition.

As you prepare for the transition, plan for travel between the sites. Note that after the remote nodes are racked

and cabled, you need serial terminal access to the nodes. Service Processor access is not be available until the nodes are configured.



Transitioning the configuration

You must follow the detailed transition procedure.

About this task

In the following steps you are directed to other procedures. You must perform the steps in each referenced procedure in the order given.

Steps

1. Plan port mapping using the steps in [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#).
2. Prepare the MetroCluster IP controllers using the steps in [Preparing the MetroCluster IP controllers](#).
3. Verify the health of the MetroCluster FC configuration.

Perform the steps in [Verifying the health of the MetroCluster FC configuration](#).

4. Gather information from the MetroCluster FC configuration.

Perform the steps in [Gathering information from the existing controller modules before the transition](#).

5. Remove Tiebreaker monitoring, if necessary.

Perform the steps in [Removing the existing configuration from the Tiebreaker or other monitoring software](#).

6. Prepare and remove the existing MetroCluster FC nodes.

Perform the steps in [Transitioning the MetroCluster FC nodes](#).

7. Connect the new MetroCluster IP nodes.

Perform the steps in [Connecting the MetroCluster IP controller modules](#).

8. Configure the new MetroCluster IP nodes and complete transition.

Perform the steps in [Configuring the new nodes and completing transition](#).

Migrating the root aggregates

After the transition is complete, migrate the existing root aggregates leftover from the MetroCluster FC configuration to new shelves in the MetroCluster IP configuration.

About this task

This task moves the root aggregates for node_A_1-FC and node_B_1-FC to disk shelves owned by the new MetroCluster IP controllers:

Steps

1. Assign pool 0 disks on the new local storage shelf to the controller that has the root being migrated (e.g., if the root of node_A_1-FC is being migrated, assign pool 0 disks on the new shelf to node_A_1-IP)

Note that the migration *removes and does not re-create the root mirror*, so pool 1 disks do not need to be assigned before issuing the migrate command

2. Set the privilege mode to advanced:

```
set priv advanced
```

3. Migrate the root aggregate:

```
system node migrate-root -node node-name -disklist disk-id1,disk-id2,diskn  
-raid-type raid-type
```

- The node-name is the node to which the root aggregate is being migrated.
- The disk-id identifies the pool 0 disks on the new shelf.
- The raid-type is normally the same as the raid-type of the existing root aggregate.
- You can use the command `job show -idjob-id-instance` to check the migration status, where job-id is the value provided when the migrate-root command is issued.

For example, if the root aggregate for node_A_1-FC consisted of three disks with raid_dp, the following command would be used to migrate root to a new shelf 11:

```
system node migrate-root -node node_A_1-IP -disklist  
3.11.0,3.11.1,3.11.2 -raid-type raid_dp
```

4. Wait until the migration operation completes and the node automatically reboots.
5. Assign pool 1 disks for the root aggregate on a new shelf directly connected to the remote cluster.
6. Mirror the migrated root aggregate.
7. Wait for the root aggregate to complete resynchronising.

You can use the storage aggregate show command to check the sync status of the aggregates.

8. Repeat these steps for the other root aggregate.

Migrating the data aggregates

Create data aggregates on the new shelves and use volume move to transfer the data volumes from the old shelves to the aggregates on the new shelves.

1. Move the data volumes to aggregates on the new controllers, one volume at a time.

[Creating an aggregate and moving volumes to the new nodes](#)

Retiring shelves moved from node_A_1-FC and node_A_2-FC

You retire the old storage shelves from the original MetroCluster FC configuration. These shelves were originally owned by node_A_1-FC and node_A_2-FC.

1. Identify the aggregates on the old shelves on cluster_B that need to be deleted.

In this example the following data aggregates are hosted by the MetroCluster FC cluster_B and need to be deleted: aggr_data_a1 and aggr_data_a2.



You need to perform the steps to identify, offline and delete the data aggregates on the shelves. The example is for one cluster only.

```
cluster_B::> aggr show

Aggregate      Size Available Used% State   #Vols  Nodes      RAID
Status
-----
-----
aggr0_node_A_1-FC
      349.0GB   16.83GB   95% online      1 node_A_1-IP
raid_dp,
mirrored,
normal
aggr0_node_A_2-IP
      349.0GB   16.83GB   95% online      1 node_A_2-IP
raid_dp,
mirrored,
normal
...
8 entries were displayed.

cluster_B::>
```

2. Check if the data aggregates have any MDV_aud volumes, and delete them prior to deleting the aggregates.

You must delete the MDV_aud volumes as they cannot be moved.

3. Take each of the aggregates offline, and then delete them:

- a. Take the aggregate offline:

```
storage aggregate offline -aggregate aggregate-name
```

The following example shows the aggregate node_B_1_aggr0 being taken offline:

```
cluster_B::> storage aggregate offline -aggregate node_B_1_aggr0

Aggregate offline successful on aggregate: node_B_1_aggr0
```

- b. Delete the aggregate:

```
storage aggregate delete -aggregate aggregate-name
```

You can destroy the plex when prompted.

The following example shows the aggregate node_B_1_aggr0 being deleted.

```
cluster_B::> storage aggregate delete -aggregate node_B_1_aggr0
Warning: Are you sure you want to destroy aggregate "node_B_1_aggr0"?
{y|n}: y
[Job 123] Job succeeded: DONE

cluster_B::>
```

4. After deleting all aggregates, power down, disconnect, and remove the shelves.
5. Repeat the above steps to retire the cluster_A shelves.

Completing transition

With the old controller modules removed, you can complete the transition process.

Step

1. Complete the transition process.

Perform the steps in [Returning the system to normal operation](#).

Disruptively transitioning when existing shelves are not supported on new controllers (ONTAP 9.8 and later)

Beginning with ONTAP 9.8, you can disruptively transition a two-node MetroCluster FC configuration and move data from the existing drive shelves even if the existing storage shelves are not supported by the new MetroCluster IP nodes.

- This procedure should only be used if the existing storage shelf models are not supported by the new MetroCluster IP platform models.
- This procedure is supported on systems running ONTAP 9.8 and later.
- This procedure is disruptive.
- This procedure applies only to a two-node MetroCluster FC configuration.

If you have a four-node MetroCluster FC configuration, see [Choosing your transition procedure](#).

- You must meet all requirements and follow all steps in the procedure.

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Requirements for transition when shelves are not supported on the new nodes

Before starting the transition process, you must make sure the configuration meets the requirements.

Before you begin

- The existing configuration must be a two-node fabric-attached or stretch MetroCluster configuration and all nodes must be running ONTAP 9.8 or later.

The new MetroCluster IP controller modules should be running the same version of ONTAP 9.8.

- The existing and new platforms must be a supported combination for transition.

[Supported platforms for nondisruptive transition](#)

- It must meet all requirements and cabling as described in [Fabric-attached MetroCluster installation and configuration](#).
- New storage shelves provided with the new controllers (node_A_1-IP, node_A_2-IP, node_B_1-IP and node_B_2-IP) must be supported by the old controllers (node_A_1-FC and node_B_1-FC).

[NetApp Hardware Universe](#)

- The old storage shelves are **not** supported by the new MetroCluster IP platform models.

[NetApp Hardware Universe](#)

- Depending on the spare disks available in the existing shelves, additional drives must be added.

This might require additional drive shelves.

You need to have additional 14 to 18 drives for each controller:

- Three pool0 drives
 - Three pool1 drives
 - Two spare drives
 - Six to ten drives for the system volume
- You must ensure that the configuration, including the new nodes, does not exceed the platform limits for the configuration, including drive count, root aggregate size capacity, etc.

This information is available for each platform model at *NetApp Hardware Universe*.

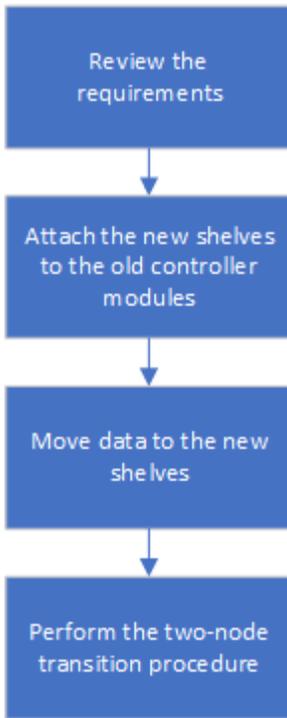
[NetApp Hardware Universe](#)

- You must have remote console access for all six nodes from either MetroCluster site or plan for travel between the sites as required by the procedure.

Workflow for disruptive transition when shelves are not supported by new controllers

If the existing shelf models are not supported by the new platform models, you must attach the new shelves to the old configuration, move data onto the new shelves, and then transition to the new configuration.

As you prepare for the transition, plan for travel between the sites. Note that after the remote nodes are racked and cabled, you need serial terminal access to the nodes. Service Processor access is not be available until the nodes are configured.



Preparing the new controller modules

You must clear the configuration and disk ownership on the new controller modules and the new storage shelves.

Steps

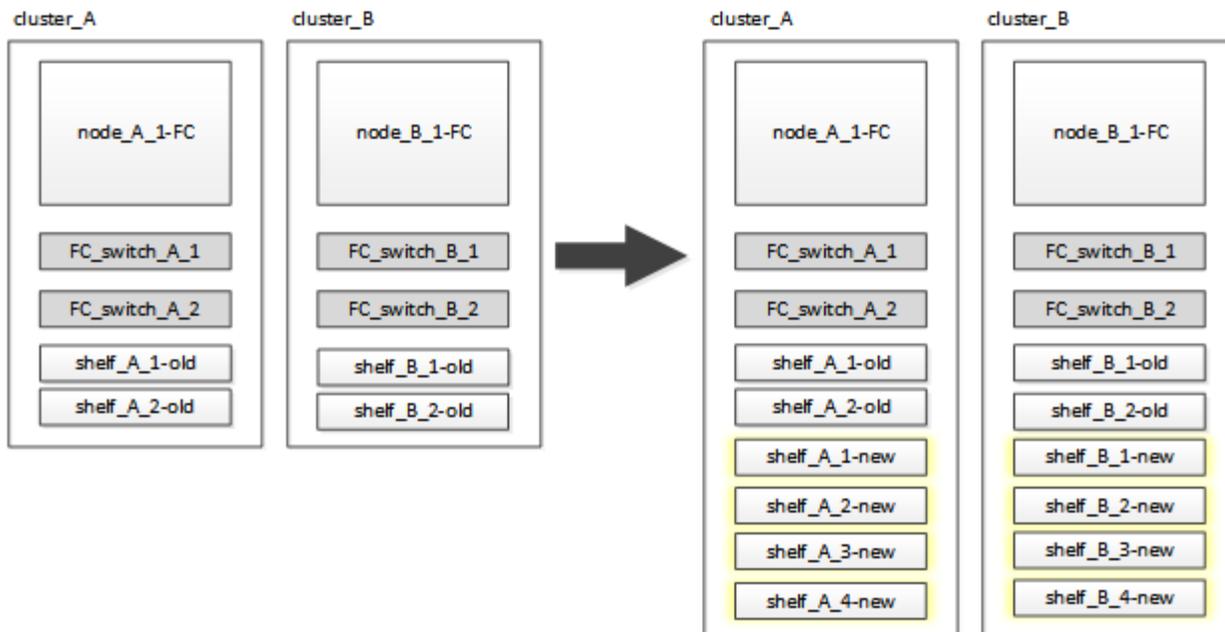
1. With the new storage shelves attached to the new MetroCluster IP controller modules, perform all the steps in [Preparing the MetroCluster IP controllers](#).
2. Disconnect the new storage shelves from the new MetroCluster IP controller modules.

Attaching the new disk shelves to the existing MetroCluster FC controllers

You must attach the new drive shelves to the existing controller modules before transitioning to a MetroCluster IP configuration.

About this task

The following illustration shows the new shelves attached to the MetroCluster FC configuration.



Steps

1. Disable disk autoassignment on node_A_1-FC and node_A_2-FC:

```
disk option modify -node node-name -autoassign off
```

This command must be issued on each node.

Disk auto-assign is disabled to avoid assignment of the shelves to be added to node_A_1-FC and node_B_1-FC. As part the transition, disks are needed for nodes node_A_1-IP and node_B_2-IP and if autoassign is allowed, disk ownership would later need to be removed before disks could be assigned to node_A_1-IP and node_B_2-IP.

2. Attach the new shelves to the existing MetroCluster FC nodes, using FC-to-SAS bridges, if necessary.

See the requirements and procedures in [Hot-adding storage to a MetroCluster FC configuration](#)

Migrate root aggregates and move data to the new disk shelves

You must move the root aggregates from the old drive shelves to the new drive shelves that will be used by the MetroCluster IP nodes.

About this task

This task is performed prior to the transition on the existing nodes (node_A_1-FC and node_B_1-FC).

Steps

1. Perform a negotiated switchover from controller node_B_1-FC:

```
metrocluster switchover
```

2. Perform the heal aggregates and heal root steps of the recovery from node_B_1-FC:

```
metrocluster heal -phase aggregates
```

```
metrocluster heal -phase root-aggregates
```

3. Boot controller node_A_1-FC:

```
boot_ontap
```

4. Assign the unowned disks on the new shelves to the appropriate pools for controller node_A_1-FC:

a. Identify the disks on the shelves:

```
disk show -shelf pool_0_shelf -fields container-type,diskpathnames
```

```
disk show -shelf pool_1_shelf -fields container-type,diskpathnames
```

b. Enter local mode so the commands are run on the local node:

```
run local
```

c. Assign the disks:

```
disk assign disk1disk2disk3disk... -p 0
```

```
disk assign disk4disk5disk6disk... -p 1
```

d. Exit local mode:

```
exit
```

5. Create a new mirrored aggregate to become the new root aggregate for controller node_A_1-FC:

a. Set the privilege mode to advanced:

```
set priv advanced
```

b. Create the aggregate:

```
aggregate create -aggregate new_aggr -disklist disk1, disk2, disk3,... -mirror  
-disklist disk4disk5, disk6,... -raidtypesame-as-existing-root -force-small  
-aggregate true aggr show -aggregate new_aggr -fields percent-snapshot-space
```

If the percent-snapshot-space value is less than 5 percent, you must increase it to a value higher than 5 percent:

```
aggr modify new_aggr -percent-snapshot-space 5
```

c. Set the privilege mode back to admin:

```
set priv admin
```

6. Confirm that the new aggregate is created properly:

```
node run -node local sysconfig -r
```

7. Create the node and cluster-level configuration backups:



When the backups are created during switchover, the cluster is aware of the switched over state on recovery. You must ensure that the backup and upload of the system configuration is successful as without this backup it is **not** possible to reform the MetroCluster configuration between clusters.

a. Create the cluster backup:

```
system configuration backup create -node local -backup-type cluster -backup  
-name cluster-backup-name
```

b. Check cluster backup creation

```
job show -id job-idstatus
```

c. Create the node backup:

```
system configuration backup create -node local -backup-type node -backup  
-name node-backup-name
```

d. Check for both cluster and node backups:

```
system configuration backup show
```

You can repeat the command until both backups are shown in the output.

8. Make copies of the backups.

The backups must be stored at a separate location because they will be lost locally when the new root volume is booted.

You can upload the backups to an FTP or HTTP server, or copy the backups using `scp` commands.

Process	Steps
Upload the backup to the FTP or HTTP server	<p>a. Upload the cluster backup:</p> <pre>system configuration backup upload -node local -backup <i>cluster-backup-name</i> -destination URL</pre> <p>b. Upload the node backup:</p> <pre>system configuration backup upload -node local -backup <i>node-backup-name</i> -destination URL</pre>

Copy the backups onto a remote server using secure copy

From the remote server use the following scp commands:

- a. Copy the cluster backup:

```
scp diagnode-mgmt-FC:/mroot/etc/backups/config/cluster-backup-name.7z .
```

- b. Copy the node backup:

```
scp diag@node-mgmt-FC:/mroot/etc/backups/config/node-backup-name.7z .
```

9. Halt node_A_1-FC:

```
halt -node local -ignore-quorum-warnings true
```

10. Boot node_A_1-FC to Maintenance mode:

```
boot_ontap maint
```

11. From Maintenance mode, make required changes to set the aggregate as root:

- a. Set the HA policy to cfo:

```
aggr options new_aggr ha_policy cfo
```

Respond “yes” when prompted to proceed.

```
Are you sure you want to proceed (y/n)?
```

- b. Set the new aggregate as root:

```
aggr options new_aggr root
```

- c. Halt to the LOADER prompt:

```
halt
```

12. Boot the controller and back up the system configuration.

The node boots in recovery mode when the new root volume is detected

- a. Boot the controller:

```
boot_ontap
```

- b. Log in and back up the configuration.

When you log in, you will see the following warning:

Warning: The correct cluster system configuration backup must be restored. If a backup from another cluster or another system state is used then the root volume will need to be recreated and NGS engaged for recovery assistance.

c. Enter advanced privilege mode:

```
set -privilege advanced
```

d. Back up the cluster configuration to a server:

```
system configuration backup download -node local -source URL of server/cluster-backup-name.7z
```

e. Back up the node configuration to a server:

```
system configuration backup download -node local -source URL of server/node-backup-name.7z
```

f. Return to admin mode:

```
set -privilege admin
```

13. Check the health of the cluster:

a. Issue the following command:

```
cluster show
```

b. Set the privilege mode to advanced:

```
set -privilege advanced
```

c. Verify the cluster configuration details:

```
cluster ring show
```

d. Return to the admin privilege level:

```
set -privilege admin
```

14. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

c. Issue the following command:

```
metrocluster check run
```

d. Display the results of the MetroCluster check:

```
metrocluster check show
```

15. Perform a switchback from controller node_B_1-FC:

```
metrocluster switchback
```

16. Verify the operation of the MetroCluster configuration:

a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

b. Perform a MetroCluster check:

```
metrocluster check run
```

c. Display the results of the MetroCluster check:

```
metrocluster check show
```

17. Add the new root volume to the Volume Location Database.

a. Set the privilege mode to advanced:

```
set -privilege advanced
```

b. Add the volume to the node:

```
volume add-other-volumes -node node_A_1-FC
```

c. Return to the admin privilege level:

```
set -privilege admin
```

18. Check that the volume is now visible and has mroot.

a. Display the aggregates:

```
storage aggregate show
```

b. Verify that the root volume has mroot:

```
storage aggregate show -fields has-mroot
```

c. Display the volumes:

```
volume show
```

19. Create a new security certificate to re-enable access to System Manager:

```
security certificate create -common-name name -type server -size 2048
```

20. Repeat the previous steps to migrate the aggregates on shelves owned by node_A_1-FC.
21. Perform a cleanup.

You must perform the following steps on both node_A_1-FC and node_B_1-FC to remove the old root volume and root aggregate.

- a. Delete the old root volume:

```
run local

vol offline old_vol0

vol destroy old_vol0

exit

volume remove-other-volume -vserver node_name -volume old_vol0
```

- b. Delete the original root aggregate:

```
aggr offline -aggregate old_aggr0_site

aggr delete -aggregate old_aggr0_site
```

22. Migrate the data volumes to aggregates on the new controllers, one volume at a time.

Refer to [Creating an aggregate and moving volumes to the new nodes](#)

23. Retire the old shelves by performing all the steps in [Retiring shelves moved from node_A_1-FC and node_A_2-FC](#).

Transitioning the configuration

You must follow the detailed transition procedure.

About this task

In the following steps you are directed to other topics. You must perform the steps in each topic in the order given.

Steps

1. Plan port mapping.

Perform all the steps in [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#).

2. Prepare the MetroCluster IP controllers.

Perform all the steps in [Preparing the MetroCluster IP controllers](#).

3. Verify the health of the MetroCluster configuration.

Perform all the steps in [Verifying the health of the MetroCluster FC configuration](#).

4. Prepare and remove the existing MetroCluster FC nodes.

Perform all the steps in [Transitioning the MetroCluster FC nodes](#).

5. Add the new MetroCluster IP nodes.

Perform all the steps in [Connecting the MetroCluster IP controller modules](#).

6. Complete the transition and initial configuration of the new MetroCluster IP nodes.

Perform all the steps in [Configuring the new nodes and completing transition](#).

Moving an FC SAN workload from MetroCluster FC to MetroCluster IP nodes

When non-disruptively transitioning from MetroCluster FC to IP nodes, you must non-disruptively move FC SAN host objects from MetroCluster FC to IP nodes.

Move an FC SAN workload from MetroCluster FC to MetroCluster IP nodes

Steps

1. Set up new FC interfaces (LIFS) on MetroCluster IP nodes:
 - a. If required, on MetroCluster IP nodes, modify FC ports to be used for client connectivity to FC target personality.

This might require a reboot of the nodes.
 - b. Create FC LIFS/interfaces on IP nodes for all SAN SVMs. Optionally, verify that the WWPNs from newly created FC LIFS are logged into the FC SAN switch
2. Update SAN zoning configuration for newly added FC LIFS on MetroCluster IP nodes.

To facilitate moving of volumes that contain LUNs actively serving data to FC SAN clients, update existing FC switch zones to allow FC SAN clients to access to LUNs on MetroCluster IP nodes.

- a. On the FC SAN switch (Cisco or Brocade), add the WWPNs of newly added FC SAN LIFS to the zone.
- b. Update, save and commit the zoning changes.
- c. From the client, check for FC initiator logins to the new SAN LIFS on the MetroCluster IP nodes:

```
sanlun lun show -p
```

At this time, the client should see and be logged in to the FC interfaces on both the MetroCluster FC and MetroCluster IP nodes. LUNs and volumes are still physically hosted on the MetroCluster FC nodes.

Because LUNs are reported only on MetroCluster FC node interfaces, the client shows only paths over FC nodes. This can be seen in the output of the `sanlun lun show -p` and `multipath -ll -d` commands.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_

[root@stemgr]# multipath -ll -d
3600a098038304646513f4f674e52774b dm-5 NETAPP ,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
|  `-- 3:0:0:4 sdk 8:160 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  `-- 2:0:0:4 sdh 8:112 active ready running

```

3. Modify the reporting nodes to add the MetroCluster IP nodes

- a. List reporting nodes for LUNs on the SVM: `lun mapping show -vserver svm-name -fields reporting-nodes -ostype linux`

Reporting nodes shown are local nodes as LUNs are physically on FC nodes A_1 and A_2.

```
cluster_A::> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux
```

vserver	path	igroup	reporting-nodes
vsa_1	/vol/vsa_1_vol1/lun_linux_2	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol1/lun_linux_3	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol2/lun_linux_4	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol3/lun_linux_7	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol4/lun_linux_8	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol4/lun_linux_9	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol6/lun_linux_12	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol6/lun_linux_13	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol7/lun_linux_14	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol8/lun_linux_17	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol9/lun_linux_18	igroup_linux	A_1,A_2
vsa_1	/vol/vsa_1_vol9/lun_linux_19	igroup_linux	A_1,A_2

12 entries were displayed.

b. Add reporting nodes to include MetroCluster IP nodes.

```
cluster_A::> lun mapping add-reporting-nodes -vserver vsa_1 -path
/vol/vsa_1_vol*/lun_linux_* -nodes B_1,B_2 -igroup igroup_linux
```

12 entries were acted on.

c. List reporting nodes and verify the presence of the new nodes:

```
cluster_A::> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux
```

vserver	path	igroup	reporting-nodes
vsa_1	/vol/vsa_1_vol1/lun_linux_2	igroup_linux	A_1,A_2,B_1,B_2
vsa_1	/vol/vsa_1_vol1/lun_linux_3	igroup_linux	A_1,A_2,B_1,B_2
vsa_1	/vol/vsa_1_vol2/lun_linux_4	igroup_linux	A_1,A_2,B_1,B_2
vsa_1	/vol/vsa_1_vol3/lun_linux_7	igroup_linux	A_1,A_2,B_1,B_2
...			

12 entries were displayed.

- d. Verify that the `sg3-utils` package is installed on the Linux host. This avoids a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.
- e. Rescan the SCSI bus on the host to discover the newly added paths: `/usr/bin/rescan-scsi-bus.sh -a`

```
[root@stemgr]# /usr/bin/rescan-scsi-bus.sh -a
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 2 0 0 0 ...
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
  Vendor: NETAPP Model: LUN C-Mode Rev: 9800
  Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.
```

- f. Display the newly added paths: `sanlun lun show -p`

Each LUN will have four paths.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
up secondary sdag host4 iscsi_lf__n4_p1_
up secondary sdah host5 iscsi_lf__n3_p1_

```

- g. On the controllers, move the volumes containing LUNs from the MetroCluster FC to the MetroCluster IP nodes.

```

cluster_A::> vol move start -vserver vsa_1 -volume vsa_1_vol1
-destination-aggregate A_1_htp_005_aggr1
[Job 1877] Job is queued: Move "vsa_1_vol1" in Vserver "vsa_1" to
aggregate "A_1_htp_005_aggr1". Use the "volume move show -vserver
vsa_1 -volume vsa_1_vol1"
command to view the status of this operation.
cluster_A::> volume move show
Vserver      Volume      State      Move Phase      Percent-Complete Time-To-
Complete
-----
-----
vsa_1      vsa_1_vol1 healthy  initializing
- -

```

- h. On the FC SAN client, display the LUN information: `sanlun lun show -p`

The FC interfaces on the MetroCluster IP nodes where the LUN now resides are updated as primary paths. If the primary path is not updated after the volume move, run `/usr/bin/rescan-iscsi-bus.sh -a` or simply wait for multipath rescanning to take place.

The primary path in the following example is the LIF on MetroCluster IP node.

```
[root@localhost ~]# sanlun lun show -p

                ONTAP Path: vsa_1:/vol/vsa_1_vol1/lun_linux_2
                  LUN: 22
                LUN Size: 2g
                  Product: cDOT
                Host Device: 3600a098038302d324e5d50305063546e
                Multipath Policy: service-time 0
                Multipath Provider: Native
-----
-----
host          vserver
path          path          /dev/      host          vserver
state        type          node       adapter      LIF
-----
-----
up           primary      sddv      host6         fc_5
up           primary      sdjx      host7         fc_6
up           secondary   sdgv      host6         fc_8
up           secondary   sdkr      host7         fc_8
```

- i. Repeat the above steps for all volumes, LUNs and FC interfaces belonging to a FC SAN host.

When completed, all LUNs for a given SVM and FC SAN host should be on MetroCluster IP nodes.

- 4. Remove the reporting nodes and re-scan paths from client.

- a. Remove the remote reporting nodes (the MetroCluster FC nodes) for the linux LUNs: `lun mapping remove-reporting-nodes -vserver vsa_1 -path * -igroup igroup_linux -remote-nodes true`

```
cluster_A::> lun mapping remove-reporting-nodes -vserver vsa_1 -path
* -igroup igroup_linux -remote-nodes true
12 entries were acted on.
```

- b. Check reporting nodes for the LUNs: `lun mapping show -vserver vsa_1 -fields reporting-nodes -ostype linux`

```

cluster_A::> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux

vserver path igroup reporting-nodes
-----
-----
vsa_1 /vol/vsa_1_vol1/lun_linux_2 igroup_linux B_1,B_2
vsa_1 /vol/vsa_1_vol1/lun_linux_3 igroup_linux B_1,B_2
vsa_1 /vol/vsa_1_vol2/lun_linux_4 igroup_linux B_1,B_2
...

12 entries were displayed.

```

c. Rescan the SCSI bus on the client: `/usr/bin/rescan-scsi-bus.sh -r`

The paths from the MetroCluster FC nodes are removed:

```

[root@stemgr]# /usr/bin/rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
sg0 changed: LU not available (PQual 1)
REM: Host: scsi2 Channel: 00 Id: 00 Lun: 00
DEL: Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
sg2 changed: LU not available (PQual 1)
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
24 device(s) removed.
[2:0:0:0]
[2:0:0:1]
...

```

d. Verify that only paths from the MetroCluster IP nodes are visible from the host: `sanlun lun show -p`

- e. If required, remove iSCSI LIFs from the MetroCluster FC nodes.

This should be done if there are no other LUNs on the nodes mapped to other clients.

Move Linux iSCSI hosts from MetroCluster FC to MetroCluster IP nodes

After you transition your MetroCluster nodes from FC to IP, you might need to move your iSCSI host connections to the new nodes.

About this task

- IPv4 interfaces are created when you set up the new iSCSI connections.
- The host commands and examples are specific to Linux operating systems.
- The MetroCluster FC nodes are called old nodes and the MetroCluster IP nodes are called new nodes.

Step 1: Set up new iSCSI connections

To move the iSCSI connections, you set up new iSCSI connections to the new nodes.

Steps

1. Create iSCSI interfaces on the new nodes and check ping connectivity from the iSCSI hosts to the new interfaces on the new nodes.

Create network interfaces

All iSCSI interfaces from the SVM should be reachable by the iSCSI host.

2. On the iSCSI host, identify the existing iSCSI connections from the host to the old node:

```
iscsiadm -m session
```

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
```

3. On the new node, verify the connections from the new node:

```
iscsi session show -vserver <svm-name>
```

```

node_A_1-new::*> iscsi session show -vserver vsa_1
  Tpgroup Initiator Initiator
Vserver Name TSIH Name ISID Alias
-----
-----
vsa_1 iscsi_lf_n1_p1_4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:01
scspr1789621001.gdl.englab.netapp.com
vsa_1 iscsi_lf_n2_p1_4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:02
scspr1789621001.gdl.englab.netapp.com
2 entries were displayed.

```

4. On the new node, list the iSCSI interfaces in ONTAP for the SVM that contains the interfaces:

```
iscsi interface show -vserver <svm-name>
```

```

sti8200mcchtp001htp_siteA::*> iscsi interface show -vserver vsa_1
  Logical Status Curr Curr
Vserver Interface TPGT Admin/Oper IP Address Node Port Enabled
-----
-----
vsa_1 iscsi_lf_n1_p1_1156 up/up 10.230.68.236 sti8200mcc-htp-001 e0g
true
vsa_1 iscsi_lf_n1_p2_1157 up/up fd20:8b1e:b255:805e::78c9 sti8200mcc-
htp-001 e0h true
vsa_1 iscsi_lf_n2_p1_1158 up/up 10.230.68.237 sti8200mcc-htp-002 e0g
true
vsa_1 iscsi_lf_n2_p2_1159 up/up fd20:8b1e:b255:805e::78ca sti8200mcc-
htp-002 e0h true
vsa_1 iscsi_lf_n3_p1_1183 up/up 10.226.43.134 sti8200mccip-htp-005 e0c
true
vsa_1 iscsi_lf_n4_p1_1188 up/up 10.226.43.142 sti8200mccip-htp-006 e0c
true
6 entries were displayed.

```

5. On the iSCSI host, run discovery on any one of the iSCSI IP addresses on the SVM to discover the new targets:

```
iscsiadm -m discovery -t sendtargets -p iscsi-ip-address
```

Discovery can be run on any IP address of the SVM, including non-iSCSI interfaces.

```
[root@scspr1789621001 ~]# iscsiadm -m discovery -t sendtargets -p
10.230.68.236:3260
10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.142:3260,1188 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.134:3260,1183 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
```

6. On the iSCSI host, login to all the discovered addresses:

```
iscsiadm -m node -L all -T node-address -p portal-address -l
```

```
[root@scspr1789621001 ~]# iscsiadm -m node -L all -T iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 -p
10.230.68.236:3260 -l
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.142,3260] (multiple)
Logging in to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.134,3260] (multiple)
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.142,3260] successful.
Login to [iface: default, target: iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6, portal:
10.226.43.134,3260] successful.
```

7. On the iSCSI host, verify the login and connections:

```
iscsiadm -m session
```

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
tcp: [3] 10.226.43.142:3260,1188 iqn.1992-
08.com.netapp:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-flash)
```

8. On the new node, verify the login and connection with the host:

```
iscsi initiator show -vserver <svm-name>
```

```
sti8200mcchtp001htp_siteA:*> iscsi initiator show -vserver vsa_1
  Tpgroup Initiator
Vserver Name          TSIH Name          ISID
Igroup Name
-----
vsa_1 iscsi_lf__n1_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:01 igroup_linux
vsa_1 iscsi_lf__n2_p1_ 4 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:02 igroup_linux
vsa_1 iscsi_lf__n3_p1_ 1 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:04 igroup_linux
vsa_1 iscsi_lf__n4_p1_ 1 iqn.2020-
01.com.netapp.englab.gdl:scspr1789621001 00:02:3d:00:00:03 igroup_linux
4 entries were displayed.
```

Result

At the end of this task, the host can see all iSCSI interfaces (on the old and new nodes) and is logged in to all those interfaces.

LUNs and volumes are still physically hosted on the old nodes. Because LUNs are reported only on the old node interfaces, the host will show only paths over the old nodes. To see this, run the `sanlun lun show -p` and `multipath -ll -d` commands on the host and examine the command outputs.

```

[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
[root@scspr1789621001 ~]# multipath -ll -d
3600a098038304646513f4f674e52774b dm-5 NETAPP ,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
| `-- 3:0:0:4 sdk 8:160 active ready running
`+- policy='service-time 0' prio=10 status=enabled
  `-- 2:0:0:4 sdh 8:112 active ready running

```

Step 2: Add the new nodes as reporting nodes

After setting up the connections to the new nodes, you add the new nodes as the reporting nodes.

Steps

1. On the new node, list reporting nodes for LUNs on the SVM:

```

lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux

```

The following reporting nodes are local nodes as LUNs are physically on old nodes node_A_1-old and node_A_2-old.

```

node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux
vserver path                                igroup      reporting-nodes
-----
-----
vsa_1    /vol/vsa_1_vol1/lun_linux_2  igroup_linux node_A_1-old,node_A_2-
old
.
.
.
vsa_1    /vol/vsa_1_vol9/lun_linux_19 igroup_linux node_A_1-old,node_A_2-
old
12 entries were displayed.

```

2. On the new node, add reporting nodes:

```

lun mapping add-reporting-nodes -vserver <svm-name> -path
/vol/vsa_1_vol*/lun_linux_* -nodes node1,node2 -igroup <igroup_name>

```

```

node_A_1-new::*> lun mapping add-reporting-nodes -vserver vsa_1 -path
/vol/vsa_1_vol*/lun_linux_* -nodes node_A_1-new,node_A_2-new
-igroup igroup_linux
12 entries were acted on.

```

3. On the new node, verify that the newly added nodes are present:

```

lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux vserver path igroup reporting-nodes

```

```

node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux vserver path igroup reporting-nodes
-----
-----
-----
vsa_1 /vol/vsa_1_voll/lun_linux_2 igroup_linux node_A_1-old,node_A_2-
old,node_A_1-new,node_A_2-new
vsa_1 /vol/vsa_1_voll/lun_linux_3 igroup_linux node_A_1-old,node_A_2-
old,node_A_1-new,node_A_2-new
.
.
.
12 entries were displayed.

```

4. The `sg3-utils` package must be installed on the Linux host. This prevents a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.

On the host, verify that the `sg3-utils` package is installed:

- For a Debian based distribution:

```
dpkg -l | grep sg3-utils
```

- For a Red Hat based distribution:

```
rpm -qa | grep sg3-utils
```

If required, install the `sg3-utils` package on the Linux host:

```
sudo apt-get install sg3-utils
```

5. On the host, rescan the SCSI bus on the host and discover the newly added paths:

```
/usr/bin/rescan-scsi-bus.sh -a
```

```

[root@stemgr]# /usr/bin/rescan-scsi-bus.sh -a
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 2 0 0 0 ...
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
  Vendor: NETAPP Model: LUN C-Mode Rev: 9800
  Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.

```

6. On the iSCSI host, list the newly added paths:

```
sanlun lun show -p
```

Four paths are shown for each LUN.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
up secondary sdag host4 iscsi_lf__n4_p1_
up secondary sdah host5 iscsi_lf__n3_p1_

```

7. On the new node, move the volume/volumes containing LUNs from the old nodes to the new nodes.

```

node_A_1-new::*> vol move start -vserver vsa_1 -volume vsa_1_vol1
-destination-aggregate sti8200mccip_htp_005_aggr1
[Job 1877] Job is queued: Move "vsa_1_vol1" in Vserver "vsa_1" to
aggregate "sti8200mccip_htp_005_aggr1". Use the "volume move show
-vserver
vsa_1 -volume vsa_1_vol1" command to view the status of this operation.
node_A_1-new::*> vol move show
Vserver   Volume           State           Move           Phase           Percent-
Complete  Time-To-Complete
-----
-----
vsa_1     vsa_1_vol1       healthy         -              initializing    -
-

```

8. When the volume move to the new nodes is complete, verify that the volume is online:

```

volume show -state

```

9. The iSCSI interfaces on the new nodes where the LUN now resides are updated as primary paths. If the primary path is not updated after the volume move, run `/usr/bin/rescan-scsi-bus.sh -a` and `multipath -v3` on the host or simply wait for multipath rescanning to take place.

In the following example, the primary path is a LIF on the new node.

```

[root@stemgr]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol16/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdag host4 iscsi_lf__n4_p1_
up secondary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
up secondary sdah host5 iscsi_lf__n3_p1_

```

Step 3: Remove reporting nodes and rescan paths

You must remove the reporting nodes and rescan the paths.

Steps

1. On the new node, remove remote reporting nodes (the new nodes) for the Linux LUNs:

```
lun mapping remove-reporting-nodes -vserver <svm-name> -path * -igroup
<igroup_name> -remote-nodes true
```

In this case, the remote nodes are old nodes.

```
node_A_1-new::*> lun mapping remove-reporting-nodes -vserver vsa_1 -path
* -igroup igroup_linux -remote-nodes true
12 entries were acted on.
```

2. On the new node, check reporting nodes for the LUNs:

```
lun mapping show -vserver <svm-name> -fields reporting-nodes -ostype
linux
```

```
node_A_1-new::*> lun mapping show -vserver vsa_1 -fields reporting-nodes
-ostype linux
vserver  path                                     igroup      reporting-nodes
-----  -
-----
vsa_1    /vol/vsa_1_vol1/lun_linux_2  igroup_linux  node_A_1-
new,node_A_2-new
vsa_1    /vol/vsa_1_vol1/lun_linux_3  igroup_linux  node_A_1-
new,node_A_2-new
vsa_1    /vol/vsa_1_vol2/lun_linux_4  group_linux   node_A_1-
new,node_A_2-new
.
.
.
12 entries were displayed.
```

3. The `sg3-utils` package must be installed on the Linux host. This prevents a `rescan-scsi-bus.sh` utility not found error when you rescan the Linux host for the newly mapped LUNs using the `rescan-scsi-bus` command.

On the host, verify that the `sg3-utils` package is installed:

- For a Debian based distribution:

```
dpkg -l | grep sg3-utils
```

- For a Red Hat based distribution:

```
rpm -qa | grep sg3-utils
```

If required, install the `sg3-utils` package on the Linux host:

```
sudo apt-get install sg3-utils
```

4. On the iSCSI host, rescan the SCSI bus:

```
/usr/bin/rescan-scsi-bus.sh -r
```

The paths that are removed are the paths from the old nodes.

```

[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
sg0 changed: LU not available (PQual 1)
REM: Host: scsi2 Channel: 00 Id: 00 Lun: 00
DEL: Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
sg2 changed: LU not available (PQual 1)
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
24 device(s) removed.
[2:0:0:0]
[2:0:0:1]
.
.
.

```

5. On the iSCSI host, verify that only paths from the new nodes are visible:

```
sanlun lun show -p
```

```
multipath -ll -d
```

Where to find additional information

You can learn more about MetroCluster configuration.

MetroCluster and miscellaneous information

Information	Subject
-------------	---------

MetroCluster IP Solution Architecture and Design, TR-4689	<ul style="list-style-type: none"> • A technical overview of the MetroCluster IP configuration and operation. • Best practices for a MetroCluster IP configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster management	<ul style="list-style-type: none"> • Understanding the MetroCluster configuration • Switchover, healing, and switchback
Disaster Recovery	<ul style="list-style-type: none"> • Disaster recovery • Forced switchover • Recovery from a multi-controller or storage failure
MetroCluster Maintenance	<ul style="list-style-type: none"> • Guidelines for maintenance in a MetroCluster FC configuration • Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches • Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration • Replacing hardware at a disaster site in a fabric-attached or stretch MetroCluster FC configuration • Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration. • Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.

MetroCluster Upgrade and Expansion	<ul style="list-style-type: none"> • Upgrading or refreshing a MetroCluster configuration • Expanding a MetroCluster configuration by adding additional nodes
MetroCluster Transition	<ul style="list-style-type: none"> • Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration
MetroCluster Upgrade, Transition, and Expansion	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
ONTAP Hardware Systems Documentation Note: The standard storage shelf maintenance procedures can be used with MetroCluster IP configurations.	<ul style="list-style-type: none"> • Hot-adding a disk shelf • Hot-removing a disk shelf
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems
ONTAP concepts	<ul style="list-style-type: none"> • How mirrored aggregates work

Upgrade, refresh, or expand the MetroCluster configuration

Start here - Choose your procedure

Start here: Choose between controller upgrade, system refresh, or expansion

Depending on the scope of the equipment upgrade, you choose a controller upgrade procedure, a system refresh procedure, or an expansion procedure.

- Controller upgrade procedures apply only to the controller modules. The controllers are replaced with a new controller model.

The storage shelf models are not upgraded.

- In switchover and switchback procedures, the MetroCluster switchover operation is used to provide nondisruptive service to clients while the controller modules on the partner cluster are upgraded.
 - In an ARL-based controller upgrade procedure, the aggregate relocation operations are used to nondisruptively move data from the old configuration to the new, upgraded configuration.
- Refresh procedures apply to the controllers and the storage shelves.

In the refresh procedures, new controllers and shelves are added to the MetroCluster configuration, creating a second DR group, and then data is nondisruptively migrated to the new nodes.

The original controllers are then retired.

- Expansion procedures add additional controllers and shelves to the MetroCluster configuration without removing any.

The procedure you use depends on the type of MetroCluster and number of existing controllers.



If SVM migration is in progress, wait until all migration processes are complete before initiating the controller upgrade or system refresh procedures. Don't start new SVM migrate operations during the upgrade or refresh process.

Upgrade type	Go to...
Controller upgrade	Choose a controller upgrade procedure
System refresh	Choose a system refresh procedure
Expansion	<ul style="list-style-type: none">• Two-node MetroCluster to four• Four-node MetroCluster FC to eight• Four-node MetroCluster IP to eight

Choose a controller upgrade procedure

The controller upgrade procedure you use depends on the platform model and type of MetroCluster configuration.

In an upgrade procedure, the controllers are replaced with a new controller model. The storage shelf models are not upgraded.

- In switchover and switchback procedures, the MetroCluster switchover operation is used to provide nondisruptive service to clients while the controller modules on the partner cluster are upgraded.
- In an ARL-based controller upgrade procedure, the aggregate relocation operations are used to nondisruptively move data from the old configuration to the new, upgraded configuration.

Supported controller upgrades

Learn about supported MetroCluster IP and FC controller upgrade combinations.

Supported MetroCluster IP controller upgrades using "system controller replace" commands

Refer to the table in [Upgrade controllers in a four-node MetroCluster IP configuration using switchover and switchback with "system controller replace" commands \(ONTAP 9.13.1 and later\)](#) for supported platforms.

All other supported MetroCluster IP controller upgrades

Find your **Source** platform from the MetroCluster controller upgrade tables in this section. If the intersection of the **Source** platform row and **Target** platform column is blank, the upgrade is not supported.

- If your platform is not listed, there is no supported controller upgrade combination.
- When you perform a controller upgrade, the old and the new platform type **must** match:
 - You can upgrade a FAS system to a FAS system, or an AFF A-Series to an AFF A-Series.
 - You cannot upgrade a FAS system to an AFF A-Series, or an AFF A-Series to an AFF C-Series.

For example, if the platform you want to upgrade is a FAS8200, you can upgrade to a FAS9000. You cannot upgrade a FAS8200 system to an AFF A700 system.

- All nodes (old and new) in the MetroCluster configuration must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.

Supported AFF and FAS MetroCluster IP controller upgrades

The following tables show the supported platform combinations for upgrading an AFF or FAS system manually in a MetroCluster IP configuration, split into two groups.

- **Group 1** shows combinations for upgrades to AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.
- **Group 2** shows combinations for upgrades to AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

The following notes apply to both groups:

- Note 1: For this upgrade use the procedure [Upgrade controllers from AFF A700/FAS9000 to AFF](#)

[A900/FAS9500 in a MetroCluster IP configuration using switchover and switchback \(ONTAP 9.10.1 or later\)](#)

- Note 2: Controller upgrades are supported on systems running ONTAP 9.13.1 or later.
- Note 3: The target platform cannot have internal drives until after the controller upgrade is complete. You can add the internal drives after the upgrade.
- Note 4: Upgrades of integrated systems (disk and controllers in the same chassis) require replacement of the controller modules while retaining the existing chassis and disks.
- Note 5: Requires IOM modules to convert the old controllers to an external SAS shelf. Refer to the [Hardware Universe](#) for supported IOM modules.

AFF and FAS combinations group 1

Review the supported combinations for upgrades to AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.

AFF and FAS		Target MetroCluster IP platform												
		AFF A150	AFF A20	FAS2750 AFF A220	FAS500f AFF C250 AFF A250	FAS50	AFF C30 AFF A30	FAS8200 AFF A300	AFF A320	FAS8300 AFF C400 AFF A400	FAS8700			
Source MetroCluster IP platform	AFF A150		Note 5											
	AFF A20													
	FAS2750		Note 5											
	AFF A220													
	FAS500f													
	AFF C250													
	AFF A250													
	FAS50													
	AFF C30													
	AFF A30													
	FAS8200													
	AFF A300													
	AFF A320													
	FAS8300													
	AFF C400													
	AFF A400													
	FAS8700													
	AFF C60													
	AFF A50													
	FAS70													
	FAS9000													
	AFF A700													
	AFF A70													
AFF C800														
AFF A800														
FAS9500														
AFF A900														
AFF C80														
FAS90														
AFF A90														
AFF A1K														

AFF and FAS combinations group 2

Review the supported combinations for upgrades to AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

AFF and FAS		Target MetroCluster IP platform										
		AFF C60	AFF A50	FAS70	FAS9000 AFF A700	AFF A70	AFF C800 AFF A800	FAS9500 AFF A900	AFF C80	FAS90	AFF A90	AFF A1K
Source MetroCluster IP platform	AFF A150											
	AFF A20											
	FAS2750											
	AFF A220											
	FAS500f											
	AFF C250											
	AFF A250											
	FAS50											
	AFF C30											
	AFF A30											
	FAS8200					Note 3		Note 2			Note 3	
	AFF A300											
	AFF A320											
	FAS8300					Note 3		Note 2	Note 3		Note 3	
	AFF C400											
	AFF A400											
	FAS8700							Note 2				
	AFF C60											
	AFF A50											
	FAS70											
FAS9000												
AFF A700					Note 3		Note 1			Note 3		
AFF A70										Note 4		
AFF C800												
AFF A800								Note 4		Note 4		
FAS9500												
AFF A900										Note 3		
AFF C80												
FAS90												
AFF A90												
AFF A1K												

Supported ASA MetroCluster IP controller upgrades

The following table shows the supported platform combinations for upgrading an ASA system manually in a MetroCluster IP configuration:

ASA		Target MetroCluster IP platform							
		ASA A150	ASA C250	ASA A250	ASA C400	ASA A400	ASA C800	ASA A800	ASA A900
Source MetroCluster IP platform	ASA A150								
	ASA C250								
	ASA A250								
	ASA C400								
	ASA A400								Note 1
	ASA C800								
	ASA A800								
ASA A900									

- Note 1: Controller upgrades are supported on systems running ONTAP 9.13.1 or later.

Supported MetroCluster FC controller upgrades

Find your **Source** platform from the MetroCluster controller upgrade tables in this section. If the intersection of the **Source** platform row and **Target** platform column is blank, the upgrade is not supported.

- If your platform is not listed, there is no supported controller upgrade combination.
- When you perform a controller upgrade, the old and the new platform type **must** match:
 - You can upgrade a FAS system to a FAS system, or an AFF A-Series to an AFF A-Series.
 - You cannot upgrade a FAS system to an AFF A-Series, or an AFF A-Series to an AFF C-Series.

For example, if the platform you want to upgrade is a FAS8200, you can upgrade to a FAS9000. You cannot upgrade a FAS8200 system to an AFF A700 system.

- All nodes (old and new) in the MetroCluster configuration must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.

Supported AFF and FAS MetroCluster FC controller upgrades

The following table shows the supported platform combinations for upgrading an AFF or FAS system in a MetroCluster FC configuration:

FAS and AFF		Target MetroCluster FC platform												
		FAS80x0	AFF80x0	FAS8200	AFF A300	FAS8300	AFF A400	FAS9000	AFF A700	FAS9500	AFF A900			
Source MetroCluster FC platform	FAS8020	Note 1		Note 1		Note 1		Note 1						
	AFF8020		Note 1		Note 1		Note 1		Note 1					
	FAS8040													
	FAS8060													
	FAS8080													
	AFF8040													
	AFF8060													
	AFF8080													
	FAS8200					Note 2		Note 2		Note 4				
	AFF A300						Note 2		Note 2		Note 4		Note 4	
	FAS8300									Note 4				
	AFF A400										Note 4		Note 4	
	FAS9000									Note 3				
	AFF A700												Note 3	
	FAS9500													
AFF A900														

- Note 1: For upgrading controllers when FCVI connections on existing FAS8020 or AFF8020 nodes use ports 1c and 1d, see the following [Knowledge base article](#).
- Note 2: Controller upgrades from AFF A300 or FAS8200 platforms using onboard ports 0e and 0f as FC-VI connections are only supported on the following systems:
 - ONTAP 9.9.1 and earlier
 - ONTAP 9.10.1P9
 - ONTAP 9.11.1P5
 - ONTAP 9.12.1GA
 - ONTAP 9.13.1 and later

For more information, review the [Public Report](#).

- Note 3: For this upgrade refer to [Upgrade controllers from AFF A700/FAS9000 to AFF A900/FAS9500 in a MetroCluster FC configuration using switchover and switchback \(ONTAP 9.10.1 or later\)](#)
- Note 4: Controller upgrades are supported on systems running ONTAP 9.13.1 or later.

Supported ASA MetroCluster FC controller upgrades

The following table shows the supported platform combinations for upgrading an ASA system in a MetroCluster FC configuration:

Source MetroCluster FC platform	Destination MetroCluster FC platform	Supported?
ASAA400	ASA A400	Yes
	ASA A900	No

Source MetroCluster FC platform	Destination MetroCluster FC platform	Supported?
ASA A900	ASA A400	No
	ASA A900	Yes (see Note 1)

- Note 1: Controller upgrades are supported on systems running ONTAP 9.14.1 or later.

Choose a procedure that uses the switchover and switchback process

After reviewing the supported upgrade combinations, choose the correct controller upgrade procedure for your configuration.

MetroCluster type	Upgrade method	ONTAP version	Procedure
IP	Upgrade with 'system controller replace' commands	9.13.1 and later	Link to procedure
FC	Upgrade with 'system controller replace' commands	9.10.1 and later	Link to procedure
FC	Manual upgrade with CLI commands (AFF A700/FAS9000 to AFF A900/FAS9500 only)	9.10.1 and later	Link to procedure
IP	Manual upgrade with CLI commands (AFF A700/FAS9000 to AFF A900/FAS9500 only)	9.10.1 and later	Link to procedure
FC	Manual upgrade with CLI commands	9.8 and later	Link to procedure
IP	Manual upgrade with CLI commands	9.8 and later	Link to procedure

Choosing a procedure using aggregate relocation

In an ARL-based controller upgrade procedure, the aggregate relocation operations are used to nondisruptively move data from the old configuration to the new, upgraded configuration.

MetroCluster type	Aggregate relocation	ONTAP version	Procedure
FC	Using "system controller replace" commands to upgrade controller models in the same chassis	9.10.1 and later	Link to procedure
FC	Using <code>system controller replace</code> commands	9.8 and later	Link to procedure
FC	Using <code>system controller replace</code> commands	9.5 through 9.7	Link to procedure
FC	Using manual ARL commands	9.8	Link to procedure
FC	Using manual ARL commands	9.7 and earlier	Link to procedure

Choosing a system refresh method

The system refresh procedure you use depends on the platform model, and type of MetroCluster configuration.

Refresh procedures apply to the controllers and the storage shelves.

In the refresh procedures, new controllers and shelves are added to the MetroCluster configuration, creating a second DR group, and then data is nondisruptively migrated to the new nodes. The original controllers are then retired.

Supported MetroCluster IP tech refresh combinations

- You must complete the tech refresh procedure before adding a new load.
- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, if you have an eight-node configuration, all eight nodes must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.
- Do not exceed any object limits of the 'lower' of the platforms in the combination. Apply the lower object limit of the two platforms.
- If the target platform limits are lower than the MetroCluster limits, you must reconfigure the MetroCluster to be at, or below, the target platform limits before you add the new nodes.
- Refer to the [Hardware universe](#) for platform limits.

Supported AFF and FAS MetroCluster IP tech refresh combinations

The following table shows the supported platform combinations for refreshing an AFF or FAS system in a MetroCluster IP configuration. The tables are split into two groups:

- **Group 1** shows combinations for AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.
- **Group 2** shows combinations for AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

The following notes apply to both groups:

- Note 1: This combination requires ONTAP 9.13.1 or later.

AFF and FAS combinations group 1

Review the system refresh combinations for AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.

AFF and FAS		Target MetroCluster IP platform									
		AFF A150	AFF A20	FAS2750 AFF A220	FAS500f AFF C250 AFF A250	FAS50	AFF C30 AFF A30	FAS8200 AFF A300	AFF A320	FAS8300 AFF C400 AFF A400	FAS8700
Source MetroCluster IP platform	AFF A150	Note 1		Note 1	Note 1					Note 1	Note 1
	AFF A20										
	FAS2750 AFF A220	Note 1		Note 1	Note 1					Note 1	Note 1
	FAS500f AFF C250 AFF A250				Note 1					Note 1	Note 1
	FAS50										
	AFF C30 AFF A30										
	FAS8200 AFF A300										
	AFF A320										
	FAS8300 AFF C400 AFF A400										
	FAS8700										
	AFF C60										
	AFF A50										
	FAS70										
	FAS9000 AFF A700										
	AFF A70										
	AFF C800 AFF A800										
	FAS9500 AFF A900										
	AFF C80										
	FAS90 AFF A90										
	AFF A1K										

AFF and FAS combinations group 2

Review the system refresh combinations for AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

AFF and FAS		Target MetroCluster IP platform									
		AFF C60	AFF A50	FAS70	FAS9000 AFF A700	AFF A70	AFF C800 AFF A800	FAS9500 AFF A900	AFF C80	FAS90 AFF A90	AFF A1K
Source MetroCluster IP platform	AFF A150				Note 1		Note 1	Note 1			
	AFF A20										
	FAS2750 AFF A220				Note 1		Note 1	Note 1			
	FAS500f AFF C250 AFF A250				Note 1		Note 1	Note 1			
	FAS50										
	AFF C30 Aff A30										
	FAS8200 AFF A300										
	AFF A320										
	FAS8300										
	AFF C400 AFF A400										
	FAS8700										
	AFF C60										
	AFF A50										
	FAS70										
	FAS9000 AFF A700										
	AFF A70										
	AFF C800 AFF A800										
	FAS9500 AFF A900										
	AFF C80										
	FAS90 AFF A90										
	AFF A1K										

Supported ASA MetroCluster IP tech refresh combinations

The following table shows the supported platform combinations for refreshing an ASA system in a MetroCluster IP configuration:

ASA		Target MetroCluster IP platform							
		ASA A150	ASA C250	ASA A250	ASA C400	ASA A400	ASA C800	ASA A800	ASA A900
Source MetroCluster IP platform	ASA A150								
	ASA C250								
	ASA A250								
	ASA C400								
	ASA A400								
	ASA C800								
	ASA A800								
	ASA A900								

Supported MetroCluster FC tech refresh combinations

- You must complete the tech refresh procedure before adding a new load.
- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, if you have an eight-node configuration, all eight nodes must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.
- Do not exceed any object limits of the 'lower' of the platforms in the combination. Apply the lower object limit of the two platforms.
- If the target platform limits are lower than the MetroCluster limits, you must reconfigure the MetroCluster to be at, or below, the target platform limits before you add the new nodes.

- Refer to the [Hardware universe](#) for platform limits.

Supported AFF and FAS MetroCluster FC tech refresh combinations

The following table shows the supported platform combinations for refreshing an AFF or FAS system in a MetroCluster FC configuration:

FAS and AFF		Destination MetroCluster FC platform							
		FAS8200	AFF A300	FAS8300	AFF A400	FAS9000	AFF A700	FAS9500	AFF A900
Source MetroCluster FC platform	FAS8200								
	AFF A300								
	FAS8300								
	AFF A400								
	FAS9000								
	AFF A700								
	FAS9500								
	AFF A900								

Supported ASA MetroCluster FC tech refresh combinations

The following table shows the supported platform combinations for refreshing an ASA system in a MetroCluster FC configuration:

Source MetroCluster FC platform	Destination MetroCluster FC platform	Supported?
ASA A400	ASA A400	Yes
	ASA A900	No
ASA A900	ASA A400	No
	ASA A900	Yes

Choose a refresh procedure

Choose the refresh procedure for your configuration from the following table:

Refresh method	Configuration type	ONTAP version	Procedure
• Method: Expand the MetroCluster configuration and then remove the old nodes	Four-node FC	9.6 and later	Link to procedure
• Method: Expand the MetroCluster configuration and then remove the old nodes	Four-node IP	9.8 and later	Link to procedure

Choose an expansion procedure

The expansion procedure you use depends on the type of MetroCluster configuration and the ONTAP version.

An expansion procedure involves adding new controllers and storage to the MetroCluster configuration. The expansion must maintain an even number of controllers on each site and the procedure you use depends on the number of nodes in the original MetroCluster configuration.

Expansion method	Configuration type	ONTAP version	Procedure
Method: Expand a two-node MetroCluster FC to four	Two-node FC	ONTAP 9 and later (platforms must be supported in ONTAP 9.2 and later)	Link to procedure
Method: Expand a four-node MetroCluster FC to eight	Four-node FC	ONTAP 9 or later	Link to procedure
Method: Expand a four-node MetroCluster IP to eight	Four-node IP	ONTAP 9.9.1 and later	Link to procedure

Upgrade controllers in four-node MetroCluster IP using switchover and switchback with "system controller replace" commands (ONTAP 9.13.1 and later)

Workflow for upgrading MetroCluster IP controllers using "system controller replace" commands (ONTAP 9.13.1 or later)

You can use this guided automated MetroCluster switchover operation to perform a non-disruptive controller upgrade on a four-node MetroCluster IP configuration running ONTAP 9.13.1 later. Other components (such as storage shelves or switches) cannot be upgraded as part of this procedure.

About this workflow

You can use this workflow to upgrade MetroCluster IP controllers running ONTAP 9.13.1 or later using switchover and switchback with `system controller replace` commands.

1

Prepare to upgrade

Review the supported upgrade combinations and requirements, and complete the required tasks to prepare your system for the controller upgrade. The automated controller upgrade process begins with a series of prechecks before you collect the configuration information and remove any existing monitoring software.

2

Upgrade your controllers

The automation operation initiates the switchover operations. After these operations complete, the operation pauses so you can prepare the network configuration of the old controllers, rack and install the new controllers, reassign the root aggregate disks, and boot the new controllers.

3

Complete the upgrade

Complete the automated controller upgrade by verifying the network reachability, repeating the upgrade tasks on the second site, and restoring monitoring configuration.

Prepare to upgrade

Supported MetroCluster IP controller upgrades using "system controller replace" commands

Before you begin the MetroCluster IP controller upgrade, you need to verify that your upgrade combination is supported.

Find your **Source** platform from the MetroCluster controller upgrade tables in this section. If the intersection of the **Source** platform row and **Target** platform column is blank, the upgrade is not supported.

Before starting the upgrade, review the following considerations to verify that your configuration is supported.

- If your platform is not listed, there is no supported controller upgrade combination.
- When you perform a controller upgrade, the old and the new platform type **must** match:
 - You can upgrade a FAS system to a FAS system, or an AFF A-Series to an AFF A-Series.
 - You cannot upgrade a FAS system to an AFF A-Series, or an AFF A-Series to an AFF C-Series.

For example, if the platform you want to upgrade is a FAS8200, you can upgrade to a FAS9000. You cannot upgrade a FAS8200 system to an AFF A700 system.

- All nodes (old and new) in the MetroCluster configuration must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.

Supported AFF and FAS MetroCluster IP controller upgrades

The following table shows the supported platform combinations for upgrading an AFF or FAS system in a MetroCluster IP configuration using "system controller replace" commands, split into two groups.

- **Group 1** shows combinations for upgrades to AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.
- **Group 2** shows combinations for upgrades to AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

The following notes apply to both groups:

- Note 1: Controller upgrades are supported on systems running ONTAP 9.13.1 or later.
- Note 2: The target platform cannot have internal drives until after the controller upgrade is complete. You can add the internal drives after the upgrade.
- Note 3: Upgrades of integrated systems (disk and controllers in the same chassis) require replacement of the controller modules while retaining the existing chassis and disks.
- Note 4: Requires IOM modules to convert the old controllers to an external SAS shelf. Refer to the [Hardware Universe](#) for the supported IOM modules.
- Note 5: Requires ONTAP 9.18.1GA or later.

AFF and FAS combinations group 1

Review the supported combinations for upgrades to AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.

FAS and AFF		Target MetroCluster IP platform									
		AFF A150	AFF A20	FAS2750 AFF A220	FAS500f AFF C250 AFF A250	FAS50	AFF C30 AFF A30	FAS8200 AFF A300	AFF A320	FAS8300 AFF C400 AFF A400	FAS8700
Source MetroCluster IP platform	AFF A150		Note 4								
	AFF A20										
	FAS2750 AFF A220		Note 4								
	FAS500f AFF C250 AFF A250						Note 3				
	FAS50										
	AFF C30 AFF A20										
	FAS8200 AFF A300										
	AFF A320										
	FAS8300 AFF C400 AFF A400										
	FAS8700										
	AFF C60										
	AFF A50										
	FAS70										
	FAS9000 AFF A700										
	AFF A70										
	AFF C800 AFF A800										
	FAS9500 AFF A900										
	AFF C80										
	FAS90 AFF A90										
	AFF A1K										

AFF and FAS combinations group 2

Review the supported combinations for upgrades to AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

FAS and AFF		Target MetroCluster IP platform										
		AFF C60	AFF A50	FAS70	FAS9000 AFF A700	AFF A70	AFF C800 AFF A800	FAS9500 AFF A900	AFF C80	FAS90	AFF A90	AFF A1K
Source MetroCluster IP platform	AFF A150											
	AFF A20											
	FAS2750											
	AFF A220											
	FAS500f											
	AFF C250											
	AFF A250											
	FAS50											
	AFF C30											
	AFF A20											
	FAS8200											
	AFF A300					Note 2		Note 1			Note 2	
	AFF A320											
	FAS8300											
	AFF C400					Note 2			Note 2		Note 2	
	AFF A400											
	FAS8700											
	AFF C60											
	AFF A50											
	FAS70									Note 5		
	FAS9000						Note 2				Note 2	
	AFF A700											
	AFF A70										Note 3 and Note 5	
	AFF C800									Note 3	Note 3	
AFF A800												
FAS9500										Note 2		
AFF A900												
AFF C80												
FAS90												
AFF A90												
AFF A1K												

Supported ASA MetroCluster IP controller upgrades

Upgrading controllers by using `system controller replace` commands on ASA systems is not supported.

Refer to [Choose an upgrade or refresh method](#) for additional procedures.

What's next?

Review the [requirements for using this upgrade procedure](#).

Requirements for using this MetroCluster IP upgrade procedure

Verify that your system meets all the requirements before performing the controller upgrade.



You must perform this procedure as written each time you upgrade your controllers.

When new platforms are released, there might be new or modified steps that must be followed. For example, when upgrading to platforms introduced in ONTAP 9.15.1 or later, you must [set the required bootarg](#) and perform other additional steps for a successful upgrade.

- You can use this procedure only for controller upgrade.

Other components in the configuration, such as storage shelves or switches, cannot be upgraded at the same time.

- The MetroCluster IP switches (switch type, vendor, and model) and the firmware version must be supported on the existing and new controllers in your upgrade configuration.

Refer to the [Hardware Universe](#) or the [IMT](#) for supported switches and firmware versions.

- The MetroCluster IP systems must be running the same ONTAP version at both sites.
- When you upgrade from systems that have more slots or ports than the new system, you need to verify that there are sufficient slots and ports on the new system.

Before you start the upgrade, refer to the [Hardware Universe](#) to verify the number of slots and ports on the new system.

- You can use this procedure to upgrade controllers in a four-node MetroCluster IP configuration using NSO based automated switchover and switchback.



Performing an upgrade using aggregate relocation (ARL) with “systems controller replace” commands is not supported for a four-node MetroCluster IP configuration.

- If it is enabled on your system, [disable end-to-end encryption](#) before performing the upgrade.
- You must use the automated NSO controller upgrade procedure to upgrade the controllers at both sites in sequence.
- This automated NSO based controller upgrade procedure gives you the capability to initiate controller replacement to a MetroCluster disaster recovery (DR) site. You can only initiate a controller replacement at one site at a time.
- To initiate a controller replacement at site A, you need to run the controller replacement start command from site B. The operation guides you to replace controllers of both the nodes at site A only. To replace the controllers at site B, you need to run the controller replacement start command from site A. A message displays identifying the site at which the controllers are being replaced.

The following example names are used in this procedure:

- cluster_A at site_A
 - Before upgrade:
 - node_A_1-old
 - node_A_2-old
 - After upgrade:
 - node_A_1-new
 - node_A_2-new
- cluster_B at site_B
 - Before upgrade:
 - node_B_1-old
 - node_B_2-old
 - After upgrade:
 - node_B_1-new
 - node_B_2-new

What's next?

[Enable console logging.](#)

Enable console logging before the MetroCluster IP upgrade

Enable console logging on your devices before performing the controller upgrade.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

What's next?

Review the information in [Set the required bootarg \(for upgrades to systems introduced in 9.15.1 or later\)](#) to confirm whether you need to set a required bootarg on the existing system.

Set the required bootarg (for MetroCluster IP upgrades to systems introduced in ONTAP 9.15.1 or later)

If you are upgrading to a system introduced in ONTAP 9.15.1 or later, you usually need to set a bootarg on the old controllers before you can start the upgrade.



If your upgrade combination is affected, setting a bootarg on all of the old controllers is required for a successful upgrade. Review the information in this section carefully to check if your upgrade combination requires that you set a bootarg, and how to set the correct bootarg for your combination.

Step 1: Determine whether you need to set a bootarg on the old controllers

Before you start the upgrade, determine if you need to set a bootarg on the old controllers using the following information:

- You **must** set a bootarg on the old controllers for supported upgrades to the following systems unless stated otherwise:
 - AFF A70, AFF A90, AFF A1K
 - FAS70, FAS90
 - AFF C80
 - AFF A50, AFF A20, AFF A30
 - AFF C30, AFF C60
 - FAS50
- You do **not** have to set any bootarg on the old controllers if your upgrade is one of the following combinations:
 - From an AFF A70 system to an AFF A90 system
 - From a FAS70 system to a FAS90 system



If your upgrade doesn't require you to set any bootarg, you can skip this task and go directly to [Prepare the system for upgrade](#).

Step 2: Determine the bootarg you need to set on the old controllers

Most affected upgrades require you to set the `hw.cxgbe.toe_keepalive_disable` bootarg on the old controllers. However, certain upgrade paths require you to set the `bootarg.siw.interop_enabled` bootarg instead.

Use the following table to determine which bootarg you need to set for your specific upgrade combination.

For this upgrade...	Set the bootarg...
From AFF A250 to AFF A30	<code>bootarg.siw.interop_enabled</code>
From AFF C250 to AFF C30	<code>bootarg.siw.interop_enabled</code>
From AFF A150 to AFF A20	<code>bootarg.siw.interop_enabled</code>
From AFF A220 to AFF A20	<code>bootarg.siw.interop_enabled</code>
All other upgrades to AFF A70, AFF A90, AFF A1K, FAS70, FAS90, AFF C80, AFF A50, AFF A30, AFF C30, AFF C60, or FAS50 systems	<code>hw.cxgbe.toe_keepalive_disable</code>
<p>Note: You don't have to set any bootarg if your upgrade is from an AFF A70 to AFF A90 system or from a FAS70 to FAS90 system.</p>	

Step 3: Set the required bootarg on the old controllers

After you've determined the bootarg required for your upgrade combination, follow the steps to set the bootarg on the old controllers.



You must set the bootarg on all of the old nodes at both sites before you start the upgrade.

Steps

1. Halt one node at both sites and allow its HA partner to perform a storage takeover of the node:

```
halt -node <node_name>
```

2. Set the required bootarg for your upgrade combination. You already determined the bootarg that you need to set by using the table in [determine which bootarg you need to set](#).

hw.cxgbe.toe_keepalive_disable

- a. At the LOADER prompt of the halted node, enter the following:

```
setenv hw.cxgbe.toe_keepalive_disable 1  
  
saveenv  
  
printenv hw.cxgbe.toe_keepalive_disable
```

bootarg.siw.interop_enabled

- a. At the LOADER prompt of the halted node, enter the following:

```
setenv bootarg.siw.interop_enabled 1  
  
saveenv  
  
printenv bootarg.siw.interop_enabled
```

3. Boot the node:

```
boot_ontap
```

4. When the node boots, perform a giveback for the node at the prompt:

```
storage failover giveback -ofnode <node_name>
```

5. Repeat the steps on all four nodes in the DR group that you are upgrading.

What's next

[Prepare the system for upgrade.](#)

Prepare the MetroCluster IP system for upgrade

To prepare for the controller upgrade, you might need to upgrade the switch reference configuration files (RCFs) depending on the old and new platform models. You then perform system prechecks, collect the configuration information, and remove existing monitoring software.

Update the MetroCluster switch RCFs before upgrading controllers

Depending on the old and new platform models, you might need to update the MetroCluster switch reference configuration files (RCFs) before you upgrade controllers.

About this task

Perform this task in the following circumstances:

- The switch RCF configuration isn't on the minimum version.
- You need to change VLAN IDs used by the back-end MetroCluster connections.

Before you begin

Determine whether you need to update the RCFs before you upgrade your controllers:

- If the switch configuration wasn't configured with the minimum supported RCF version, you need to update the RCFs before you upgrade your controllers:

Switch model	Required RCF version
Cisco 3132Q-V	1.7 or later
Cisco 3232C	1.7 or later
Broadcom BES-53248	1.3 or later
NVIDIA SN2100	2.0 or later

- If both of your old and new platform models are in the following list, you do **not** need to update the VLAN ID before you upgrade controllers:
 - FAS8200 or AFF A300
 - AFF A320
 - FAS9000 or AFF A700
 - AFF A800, AFF C800, ASA A800, or ASA C800

If either of your old or new platform models are not listed above, you must confirm that the MetroCluster interfaces are using a supported VLAN ID. Supported VLAN IDs for the MetroCluster interfaces are: 10, 20, or in the range of 101 to 4096.



- If the VLAN ID is not 10, 20, or in the range of 101 to 4096, you must upgrade the switch RCF before you upgrade controllers.
- The local cluster connections can use any VLAN, they don't need to be in the given range.
- The new RCF that you are upgrading to must use the VLANs 10, 20, or be in the range 101 to 4096. Don't change the VLAN for the local cluster unless it is required.

Steps

1. Prepare the IP switches for the application of the new RCFs.

Follow the steps in the section for your switch vendor:



You should update the switches in the following order: switch_A_1, switch_B_1, switch_A_2, switch_B_2.

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

2. Download and install the RCFs.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom RCFs](#)
- [Download and install the Cisco IP RCFs](#)
- [Download and install the NVIDIA IP RCFs](#)

Start the controller replacement operation

When you start the automated controller replacement operation, the operation issues a series of prechecks and then pauses so that you can manually collect the configuration related information.

About this task

Before the MetroCluster checks start, if ONTAP Mediator is installed, it is automatically detected and removed. To confirm the removal, you are prompted to enter a username and password. When you complete the upgrade, if prechecks fail, or you choose not to proceed with the upgrade, you must manually [re-configure the ONTAP Mediator service](#).

At any stage during the upgrade, you can run the `system controller replace show` or `system controller replace show-details` command from site A to check the status. If the commands return a blank output, wait for a few minutes and rerun the command.

Steps

1. Run the following command from site A to replace the controllers at site B:

```
system controller replace start -nso true
```



- If you don't use the `-nso true` parameter in the command, the controller upgrade procedure chooses NSO based automated switchover and switchback as the default procedure on MetroCluster IP systems.
- If you're repeating the procedure at one site, after already replacing the controllers at the other site, an error occurs due to a mismatch between the nodes at each site. This is the expected behavior when there are different platform models on both sites.

If only the mismatch error is returned, you can use the `-skip-metrocluster-check true` option with the `system controller replace start` command to skip the MetroCluster checks.

The automated operation executes the checks. If no issues are found, the operation pauses so you can manually collect the configuration related information.

The current source system and all compatible target systems are displayed. If you have replaced the source controller with a controller that has a different ONTAP version or a non-compatible platform, the automation operation halts and reports an error after the new nodes boot. To bring the cluster back to a healthy state, follow the manual recovery procedure.

The `system controller replace start` command might report the following precheck error:

```
Cluster-A::*>system controller replace show
Node          Status          Error-Action
-----
Node-A-1     Failed          MetroCluster check failed. Reason : MCC check
showed errors in component aggregates
```

Check if this error occurred because you have unmirrored aggregates or due to another aggregate issue. Verify that all mirrored aggregates are healthy and not degraded or mirror-degraded. If this error is due to unmirrored aggregates only, you can override this error by selecting the `-skip-metrocluster-check true` option on the `system controller replace start` command. If remote storage is accessible, the unmirrored aggregates come online after switchover. If the remote storage link fails, the unmirrored aggregates fail to come online.

2. Manually collect the configuration information by logging in at site B and following the commands listed in the console message under the `system controller replace show` or `system controller replace show-details` command.

Gather information before the upgrade

Before upgrading, if the root volume is encrypted, you must gather the backup key and other information to boot the new controllers with the old encrypted root volumes.

About this task

This task is performed on the existing MetroCluster IP configuration.

Steps

1. Label the cables for the existing controllers, so you can easily identify the cables when setting up the new controllers.
2. Display the commands to capture the backup key and other information:

```
system controller replace show
```

Run the commands listed under the `show` command from the partner cluster.

The `show` command output displays three tables containing the MetroCluster interface IPs, system IDs, and system UUIDs. This information is required later in the procedure to set the bootargs when you boot the new node.

3. Gather the system IDs of the nodes in the MetroCluster configuration:

```
metrocluster node show -fields node-systemid,dr-partner-systemid
```

During the upgrade procedure, you will replace these old system IDs with the system IDs of the new controller modules.

In this example for a four-node MetroCluster IP configuration, the following old system IDs are retrieved:

- `node_A_1-old: 4068741258`
- `node_A_2-old: 4068741260`
- `node_B_1-old: 4068741254`

- node_B_2-old: 4068741256

```
metrocluster-siteA::> metrocluster node show -fields node-systemid,ha-
partner-systemid,dr-partner-systemid,dr-auxiliary-systemid
dr-group-id          cluster          node          node-systemid
ha-partner-systemid  dr-partner-systemid  dr-auxiliary-systemid
-----
-----
1                    Cluster_A          Node_A_1-old  4068741258
4068741260          4068741256
1                    Cluster_A          Node_A_2-old  4068741260
4068741258          4068741254
1                    Cluster_B          Node_B_1-old  4068741254
4068741256          4068741258
1                    Cluster_B          Node_B_2-old  4068741256
4068741254          4068741260
4 entries were displayed.
```

In this example for a two-node MetroCluster IP configuration, the following old system IDs are retrieved:

- node_A_1: 4068741258
- node_B_1: 4068741254

```
metrocluster node show -fields node-systemid,dr-partner-systemid
dr-group-id cluster      node          node-systemid dr-partner-systemid
-----
1          Cluster_A  Node_A_1-old  4068741258    4068741254
1          Cluster_B  node_B_1-old  -              -
2 entries were displayed.
```

4. Gather port and LIF information for each old node.

You should gather the output of the following commands for each node:

- network interface show -role cluster,node-mgmt
- network port show -node <node-name> -type physical
- network port vlan show -node <node-name>
- network port ifgrp show -node <node-name> -instance
- network port broadcast-domain show
- network port reachability show -detail
- network ipspace show

- volume show
- storage aggregate show
- system node run -node <node-name> sysconfig -a
- aggr show -r
- disk show
- system node run <node-name> disk show
- vol show -fields type
- vol show -fields type , space-guarantee
- vservers fcp initiator show
- storage disk show
- metrocluster configuration-settings interface show

5. If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

You should gather the output of the following commands:

- fcp adapter show -instance
- fcp interface show -instance
- iscsi interface show
- ucaadmin show

6. If the root volume is encrypted, collect and save the passphrase used for key-manager:

```
security key-manager backup show
```

7. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.

For additional information, see [Back up onboard key management information manually](#).

a. If Onboard Key Manager is configured:

```
security key-manager onboard show-backup
```

You will need the passphrase later in the upgrade procedure.

b. If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
```

```
security key-manager key query
```

8. After you finish collecting the configuration information, resume the operation:

```
system controller replace resume
```

Remove the existing configuration from Tiebreaker or other monitoring software

Before you start the upgrade, remove the existing configuration from the Tiebreaker or other monitoring software.

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the Tiebreaker or other software prior to replacing the old controller.

Steps

1. [Remove the existing MetroCluster configuration](#) from the Tiebreaker software.
2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

What's next?

[Prepare the network configuration of the old controllers.](#)

Upgrade your controllers

Prepare the network configuration of the old MetroCluster IP controllers

After you gather information and resume the operation, the automation operation proceeds with the switchover.

Before you begin

Before initiating switchover, the automation operation pauses so you can manually verify that all LIFs are “up” at site B. If necessary, bring any LIFs that are “down” to “up” and resume the automation operation by using the `system controller replace resume` command.

The automation operation initiates the switchover operations. After these operations complete, the operation pauses at **paused for user intervention** so you can rack and install the controllers, boot the partner controllers, and reassign the root aggregate disks to the new controller module from flash backup using the `sysids` gathered earlier.

About this task

- This task must be performed on each of the old nodes.
- You use the information gathered in [Gather information before the upgrade](#).

Steps

1. Boot the old nodes and then log in to the nodes:

```
boot_ontap
```

2. If the system you are upgrading to uses **shared cluster/HA ports**, verify that the MetroCluster IP interfaces are using supported IP addresses.

Use the following information to determine whether the new system uses shared cluster/HA ports:

Shared cluster/HA ports

The systems listed in the following table use shared cluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A20• AFF A30• AFF C30• AFF A50• AFF C60• AFF C80• AFF A70• AFF A90• AFF A1K	<ul style="list-style-type: none">• FAS50• FAS70• FAS90

Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A150, ASA A150• AFF A220• AFF C250, ASA C250• AFF A250, ASA A250• AFF A300• AFF A320• AFF C400, ASA C400• AFF A400, ASA A400• AFF A700• AFF C800, ASA C800• AFF A800, ASA A800• AFF A900, ASA A900	<ul style="list-style-type: none">• FAS2750• FAS500f• FAS8200• FAS8300• FAS8700• FAS9000• FAS9500

a. Verify the IP addresses of the MetroCluster interfaces on the old controllers:

```
metrocluster configuration-settings interface show
```

b. If the MetroCluster interfaces are using 169.254.17.x or 169.254.18.x IP addresses, refer to [the Knowledge Base article "How to modify the properties of a MetroCluster IP interface"](#) to modify the interface IP addresses before you proceed with the upgrade.



Upgrading to any system that uses **shared cluster/HA ports** isn't supported if the MetroCluster interfaces are configured with 169.254.17.x or 169.254.18.x IP addresses.

3. Modify the intercluster LIFs on the old controllers to use a different home port than the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers.



This step is required for a successful upgrade.

The intercluster LIFs on the old controllers must use a different home port than the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers. For example, when you upgrade to AFF A90 controllers, the HA interconnect ports are e1a and e7a, and the MetroCluster IP DR interconnect ports are e2b and e3b. You must move the intercluster LIFs on the old controllers if they are hosted on ports e1a, e7a, e2b, or e3b.

For port distribution and allocation on the new nodes, refer to the [Hardware Universe](#).

- a. On the old controllers, view the intercluster LIFs:

```
network interface show -role intercluster
```

Take one of the following actions depending on whether the intercluster LIFs on the old controllers use the same ports as the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers.

If the intercluster LIFs...	Go to...
Use the same home port	Substep b
Use a different home port	Step 4

- b. Modify the intercluster LIFs to use a different home port:

```
network interface modify -vserver <vserver> -lif <intercluster_lif> -home
-port <port-not-used-for-ha-interconnect-or-mcc-ip-dr-interconnect-on-new-
nodes>
```

- c. Verify that all intercluster LIFs are on their new home ports:

```
network interface show -role intercluster -is-home false
```

The command output should be empty, indicating that all intercluster LIFs are on their respective home ports.

- d. Revert any LIFs that are not on their home ports:

```
network interface revert -lif <intercluster_lif>
```

Repeat the command for each intercluster LIF that is not on the home port.

4. Assign the home port of all data LIFs on the old controller to a common port that is the same on both the old and new controller modules.



If the old and new controllers don't have a common port, you don't need to modify the data LIFs. Skip this step and go directly to [Step 5](#).

a. Display the LIFs:

```
network interface show
```

All data LIFs including SAN and NAS will be admin “up” and operationally “down” since those are up at switchover site (cluster_A).

b. Review the output to find a common physical network port that is the same on both the old and new controllers that is not used as a cluster port.

For example, “e0d” is a physical port on old controllers and is also present on new controllers. “e0d” is not used as a cluster port or otherwise on the new controllers.

Refer to the [Hardware Universe](#) for the port usage of each platform model.

c. Modify all data LIFS to use the common port as the home port:

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-port <port-id>
```

In the following example, this is “e0d”.

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

5. Modify broadcast domains to remove the VLAN and physical ports that need to be deleted:

```
broadcast-domain remove-ports -broadcast-domain <broadcast-domain-name>-ports <node-name:port-id>
```

Repeat this step for all VLAN and physical ports.

6. Remove any VLAN ports using cluster ports as member ports and interface groups using cluster ports as member ports.

a. Delete VLAN ports:

```
network port vlan delete -node <node-name> -vlan-name <portid-vlandid>
```

For example:

```
network port vlan delete -node node1 -vlan-name elc-80
```

b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node <node-name> -ifgrp <interface-group-name> -port <portid>
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp a1a -port e0d
```

c. Remove VLAN and interface group ports from the broadcast domain:

```
network port broadcast-domain remove-ports -ip-space <ip-space> -broadcast  
-domain <broadcast-domain-name>-ports  
<nodename:portname,nodename:portname>, ..
```

d. Modify interface group ports to use other physical ports as member as needed:

```
ifgrp add-port -node <node-name> -ifgrp <interface-group-name> -port <port-  
id>
```

7. Halt the nodes:

```
halt -inhibit-takeover true -node <node-name>
```

This step must be performed on both nodes.

8. Verify that the nodes are at the `LOADER` prompt and collect and preserve the current environment variables.

9. Gather the bootarg values:

```
printenv
```

10. Power off the nodes and shelves at the site where the controller is being upgraded.

What's next?

[Set up and netboot the new controllers.](#)

Set up and netboot the new MetroCluster IP controllers

Set up the new controllers before you netboot the controllers to confirm the new nodes are running the same version of ONTAP as the original nodes.

Set up the new controllers

You must rack and cable the new controllers.

Steps

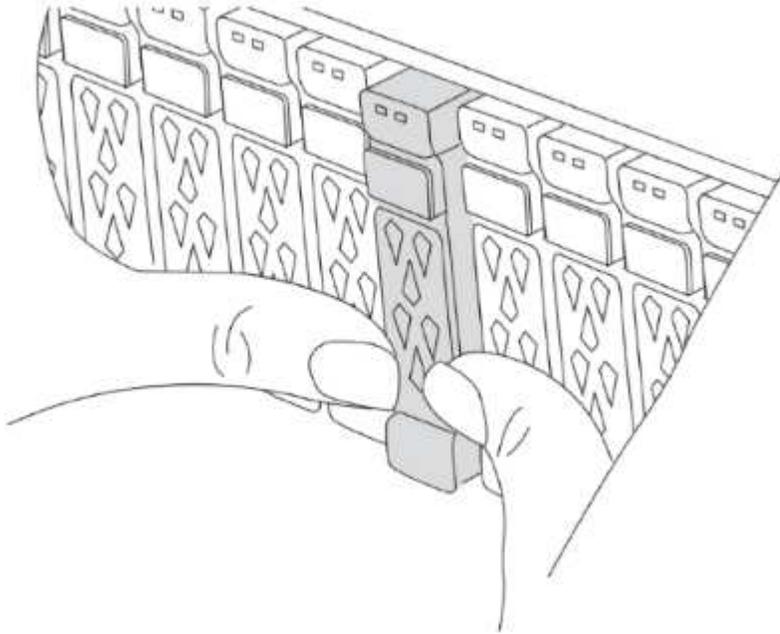
1. Plan out the positioning of the new controller modules and storage shelves as needed.

The rack space depends on the platform model of the controller modules, the switch types, and the number of storage shelves in your configuration.

2. Properly ground yourself.

3. If your upgrade requires replacement of the controller modules, for example, upgrading from an AFF A800 to an AFF A90 system or from an AFF C800 to an AFF C80 system, you must remove the controller module from the chassis when you replace the controller module. For all other upgrades, skip to [Step 4](#).

On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This confirms that the drives are firmly seated against the chassis midplane.



4. Install the controller modules.

The installation steps you follow depend on whether your upgrade requires replacement of the controller modules, or if IOM modules are required to convert the old controllers to an external shelf.

If you are upgrading...	Follow the steps for ...
<ul style="list-style-type: none"> • An AFF A150 to an AFF A20 system • An AFF A220 to an AFF A20 system 	Controller to external shelf conversion
<ul style="list-style-type: none"> • An AFF A800 to an AFF A90 system • An AFF C800 to an AFF C80 system • An AFF A250 to an AFF A30 system • An AFF C250 to an AFF C30 system • An AFF A70 to an AFF A90 system 	Controller module replacement
Any other controller upgrade combinations	All other upgrades

Controller to external shelf conversion

If your original MetroCluster IP controllers are AFF A150 or AFF A220 models, you can convert the AFF A150 or AFF A220 HA pair to a DS224C drive shelf and then attach it to the new nodes.

For example, when upgrading from an AFF A150 or AFF A220 system to an AFF A20 system, you can convert the AFF A150 or AFF A220 HA pair to a DS224C shelf by swapping the AFF A150 or AFF A220 controller modules with IOM12 modules.

Steps

1. Replace the controller modules in the node you are converting with IOM12 shelf modules.

[Hardware Universe](#)

2. Set the drive shelf ID.

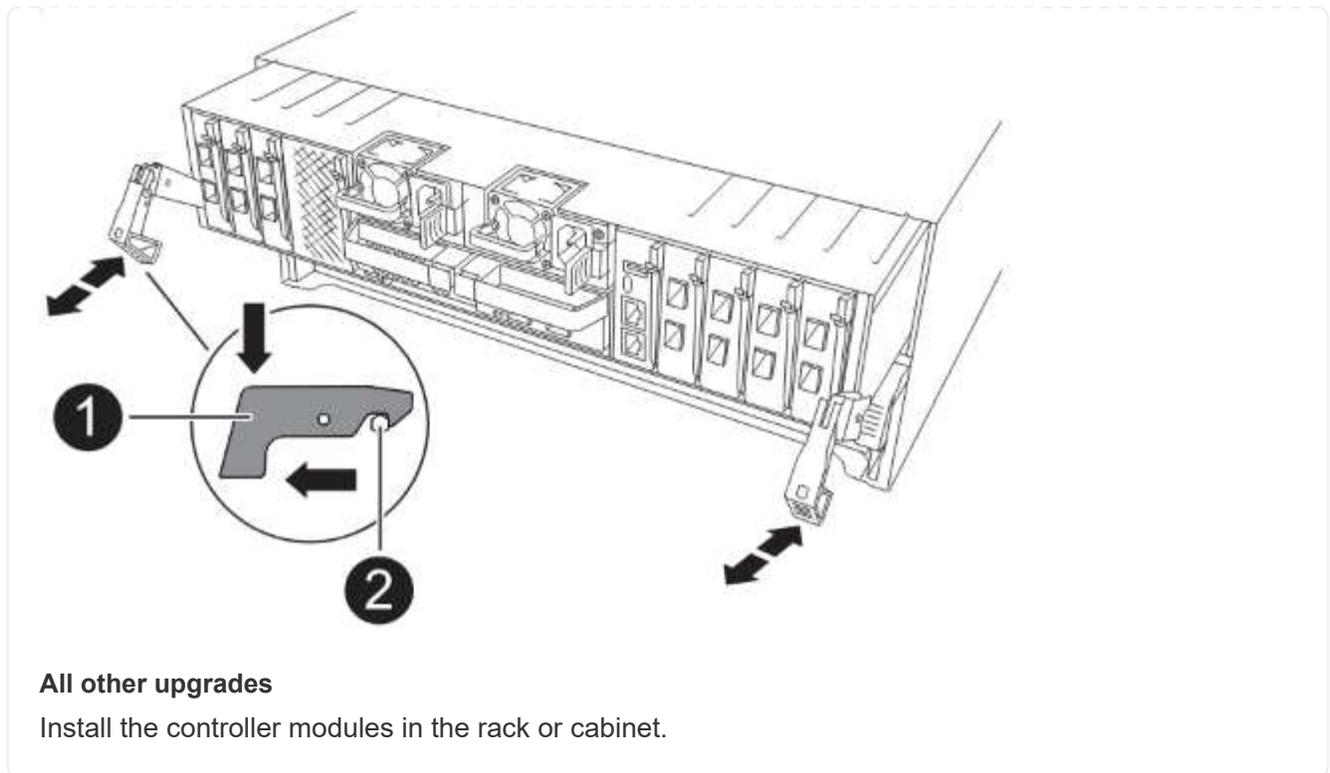
Each drive shelf, including the chassis, requires a unique ID.

3. Reset other drive shelf IDs as needed.
4. Power off the shelves.
5. Cable the converted drive shelf to a SAS port on the new system, and, if you are using out-of-band ACP cabling, to the ACP port on the new node.
6. Turn on the power to the converted drive shelf and any other drive shelves attached to the new nodes.
7. Turn on the power to the new nodes, and then interrupt the boot process on each node by pressing Ctrl-C to access the boot environment prompt.

Controller module replacement

Installing the new controllers separately is not applicable for upgrades of integrated systems with disks and controllers in the same chassis, for example, from an AFF A800 system to an AFF A90 system. You must swap the new controller modules and I/O cards after powering off the old controllers, as shown in the image below.

The following example image is for representation only, the controller modules and I/O cards can vary between systems.



5. Cable the controllers' power, serial console, and management connections as described in [Cable the MetroCluster IP switches](#)

Do not connect any other cables that were disconnected from old controllers at this time.

[ONTAP Hardware Systems Documentation](#)

6. Power up the new nodes and press Ctrl-C when prompted to display the `LOADER` prompt.

Netboot the new controllers

After you install the new nodes, you need to netboot to ensure the new nodes are running the same version of ONTAP as the original nodes. The term netboot means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

This task is performed on each of the new controller modules.

Steps

1. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `ontap-version_image.tgz` file on a web-accessible directory.
3. Go to the web-accessible directory and verify that the files you need are available.

Your directory listing should contain a netboot folder with a kernel file: `ontap-version_image.tgz`

You do not need to extract the `ontap-version_image.tgz` file.

4. At the `LOADER` prompt, configure the netboot connection for a management LIF:

- If IP addressing is DHCP, configure the automatic connection:

```
ifconfig e0M -auto
```

- If IP addressing is static, configure the manual connection:

```
ifconfig e0M -addr=ip_addr -mask=netmask -gw=gateway
```

5. Perform the netboot.

```
netboot http://web_server_ip/path_to_web-accessible_directory/ontap-  
version_image.tgz
```

6. From the boot menu, select option **(7) Install new software first** to download and install the new software image to the boot device.

Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". It applies to nondisruptive upgrades of software, not to upgrades of controllers.

7. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL of the image file: `http://web_server_ip/path_to_web-accessible_directory/ontap-version_image.tgz`

Enter username/password if applicable, or press Enter to continue.

8. Enter `n` to skip the backup recovery when you see a prompt similar to the following:

Do you want to restore the backup configuration now? {y|n} n

9. Reboot by entering `y` when you see a prompt similar to the following:

The node must be rebooted to start using the newly installed software.
Do you want to reboot now? {y|n} y



You must reboot the node in order to use the newly installed software.

Clear the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing configuration.

Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

What's next?

[Restore the HBA configuration and set the HA state.](#)

Restore the HBA configuration and set the HA state of the MetroCluster IP controller and chassis

Configure the HBA cards in the controller module and verify and set the HA state of the controller and chassis.

Restore the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site.

Steps

1. In Maintenance mode, configure the settings for any HBAs in the system:
 - a. Check the current settings of the ports: `ucadmin show`
 - b. Update the port settings as needed.

If you have this type of HBA and desired mode...	Use this command...
CNA FC	<pre>ucadmin modify -m fc -t initiator <adapter-name></pre>

CNA Ethernet	<code>ucadmin modify -mode cna <adapter-name></code>
FC target	<code>fcadmin config -t target <adapter-name></code>
FC initiator	<code>fcadmin config -t initiator <adapter-name></code>

2. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the `LOADER` prompt.

3. Boot the node back into Maintenance mode to apply the configuration changes:

```
boot_ontap maint
```

4. Verify the changes:

If you have this type of HBA...	Use this command...
CNA	<code>ucadmin show</code>
FC	<code>fcadmin show</code>

Set the HA state on the new controllers and chassis

You must verify the HA state of the controllers and chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be `mccip`.

2. If the displayed system state of the controller or chassis is not correct, set the HA state:

```
ha-config modify controller mccip
```

```
ha-config modify chassis mccip
```

3. Verify and modify the Ethernet ports connected to NS224 shelves or storage switches.

a. Verify the Ethernet ports connected to NS224 shelves or storage switches:

```
storage port show
```

- b. Set all Ethernet ports connected to Ethernet shelves or storage switches, including shared switches for storage and cluster, to `storage` mode:

```
storage port modify -p <port> -m storage
```

Example:

```
*> storage port modify -p e5b -m storage
Changing NVMe-oF port e5b to storage mode
```



This must be set on all affected ports for a successful upgrade.

Disks from the shelves attached to the Ethernet ports are reported in the `sysconfig -v` output.

Refer to the [Hardware Universe](#) for information on the storage ports for the system you are upgrading to.

- c. Verify that `storage` mode is set and confirm that the ports are in the online state:

```
storage port show
```

4. Halt the node: `halt`

The node should stop at the `LOADER>` prompt.

5. On each node, check the system date, time, and time zone: `show date`
6. If necessary, set the date in UTC or GMT: `set date <mm/dd/yyyy>`
7. Check the time by using the following command at the boot environment prompt: `show time`
8. If necessary, set the time in UTC or GMT: `set time <hh:mm:ss>`
9. Save the settings: `saveenv`
10. Gather environment variables: `printenv`

Remove internal drives from the chassis on the new controller

When you upgrade from a system that only has external drives to a system that has external and internal drives (disks and controllers in the same chassis), you need to remove or unseat all internal drives from the new system until you have completed the upgrade.



This task is mandatory for a successful controller upgrade on affected systems.

To determine if your upgrade combination is affected, refer to the table in [Supported MetroCluster IP controller upgrades using "system controller replace" commands](#). If your upgrade combination is marked with **Note 2**, you must remove or unseat the internal drives from the new system.

After you complete this task, no internal drives should be accessible. You'll add the drives to the new controller later in the procedure.

What's next?

[Update the switch RCF files and set the MetroCluster IP bootarg values.](#)

Update the switch RCFs and set the MetroCluster IP bootarg values

Update the switch reference configuration files (RCFs) for the new platforms and set the MetroCluster IP bootarg values on the controller modules.

Update the switch RCFs to accommodate the new platforms

You must update the switches to a configuration that supports the new platform models.

About this task

You perform this task at the site containing the controllers that are currently being upgraded. In the examples shown in this procedure we are upgrading site_B first.

The switches at site_A will be upgraded when the controllers on site_A are upgraded.

Steps

1. Prepare the IP switches for the application of the new RCFs.

Follow the steps in the section for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

2. Download and install the RCFs.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom RCFs](#)
- [Download and install the Cisco IP RCFs](#)
- [Download and install the NVIDIA IP RCFs](#)

Set the MetroCluster IP bootarg variables

Certain MetroCluster IP bootarg values must be configured on the new controller modules. The values must match those configured on the old controller modules.

About this task

- You need the UUIDs and system IDs identified earlier in the upgrade procedure in [Gather information before the upgrade](#).
- Depending on your platform model, you can specify the VLAN ID using the `-vlan-id` parameter. The following platforms do not support the `-vlan-id` parameter:
 - FAS8200 and AFF A300
 - AFF A320
 - FAS9000 and AFF A700
 - AFF C800, ASA C800, AFF A800, and ASA A800

All other platforms support the `-vlan-id` parameter.

- The MetroCluster bootarg values you set depend on whether your new system uses shared cluster/HA ports or shared MetroCluster/HA ports.

Shared cluster/HA ports

The systems listed in the following table use shared cluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none"> • AFF A20 • AFF A30 • AFF C30 • AFF A50 • AFF C60 • AFF C80 • AFF A70 • AFF A90 • AFF A1K 	<ul style="list-style-type: none"> • FAS50 • FAS70 • FAS90

Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none"> • AFF A150, ASA A150 • AFF A220 • AFF C250, ASA C250 • AFF A250, ASA A250 • AFF A300 • AFF A320 • AFF C400, ASA C400 • AFF A400, ASA A400 • AFF A700 • AFF C800, ASA C800 • AFF A800, ASA A800 • AFF A900, ASA A900 	<ul style="list-style-type: none"> • FAS2750 • FAS500f • FAS8200 • FAS8300 • FAS8700 • FAS9000 • FAS9500

Steps

1. At the `LOADER>` prompt, set the following bootargs on the new nodes at `site_B`:

The steps you follow depends on the ports used by the new platform model.

Systems that use shared cluster/HA ports

- a. Set the following bootargs:

```
setenv bootarg.mcc.port_a_ip_config <local-IP-address/local-IP-  
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id>
```

```
setenv bootarg.mcc.port_b_ip_config <local-IP-address/local-IP-  
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id>
```



If the interfaces are using a default VLAN ID, the `vlan-id` parameter is not required.

The following example sets the values for node_B_1-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12,130
```

The following example sets the values for node_B_2-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,0,172.17.26.12,172.17.26.13,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,0,172.17.27.12,172.17.27.13,130
```

The following example sets the values for node_B_1-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12
```

The following example sets the values for node_B_2-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,0,172.17.26.12,172.17.26.13  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,0,172.17.27.12,172.17.27.13
```

Systems that use shared MetroCluster/HA ports

- a. Set the following bootargs:

```
setenv bootarg.mcc.port_a_ip_config <local-IP-address/local-IP-  
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-  
address,vlan-id>
```

```
setenv bootarg.mcc.port_b_ip_config <local-IP-address/local-IP-  
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-  
address,vlan-id>
```



If the interfaces are using a default VLAN ID, the `vlan-id` parameter is not required.

The following example sets the values for node_B_1-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12,130
```

The following example sets the values for node_B_2-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13,130
```

The following example sets the values for node_B_1-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12
```

The following example sets the values for node_B_2-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13
setenv bootarg.mcc.port_b_ip_config
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13
```

2. At the new nodes' LOADER prompt, set the UUIDs:

```
setenv bootarg.mgwd.partner_cluster_uuid <partner-cluster-UUID>
setenv bootarg.mgwd.cluster_uuid <local-cluster-UUID>
setenv bootarg.mcc.pri_partner_uuid <DR-partner-node-UUID>
setenv bootarg.mcc.aux_partner_uuid <DR-aux-partner-node-UUID>
setenv bootarg.mcc.iscsi.node_uuid <local-node-UUID>
```

a. Set the UUIDs on node_B_1-new.

The following example shows the commands for setting the UUIDs on node_B_1-new:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.pri_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.aux_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f
setenv bootarg.mcc.iscsi.node_uuid f03cb63c-9a7e-11e7-b68b-
00a098908039
```

b. Set the UUIDs on node_B_2-new:

The following example shows the commands for setting the UUIDs on node_B_2-new:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.pri_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f
setenv bootarg.mcc.aux_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.iscsi.node_uuid aa9a7a7a-9a81-11e7-a4e9-
00a098908c35
```

- Determine whether the original systems were configured for Advanced Drive Partitioning (ADP) by running the following command from the site that is up:

```
disk show
```

The "container type" column displays "shared" in the `disk show` output if ADP is configured. If "container type" has any other value, ADP is not configured on the system. The following example output shows a system configured with ADP:

```
::> disk show
```

Disk Owner	Usable Size	Shelf	Bay	Disk Type	Container Type	Container Name
1.11.0 node_A_1	894.0GB	11	0	SSD	shared	testaggr
1.11.1 node_A_1	894.0GB	11	1	SSD	shared	testaggr
1.11.2 node_A_1	894.0GB	11	2	SSD	shared	testaggr

Info: This cluster has partitioned disks. To get a complete list of spare disk capacity use "storage aggregate show-spare-disks".

- If the original systems were configured with partitioned disks for ADP, enable it at the `LOADER` prompt for each replacement node:

```
setenv bootarg.mcc.adp_enabled true
```

- Set the following variables:

```
setenv bootarg.mcc.local_config_id <original-sys-id>
```

```
setenv bootarg.mcc.dr_partner <dr-partner-sys-id>
```



The `setenv bootarg.mcc.local_config_id` variable must be set to the sys-id of the **original** controller module, `node_B_1-old`.

- Set the variables on `node_B_1-new`.

The following example shows the commands for setting the values on `node_B_1-new`:

```
setenv bootarg.mcc.local_config_id 537403322
setenv bootarg.mcc.dr_partner 537403324
```

b. Set the variables on node_B_2-new.

The following example shows the commands for setting the values on node_B_2-new:

```
setenv bootarg.mcc.local_config_id 537403321
setenv bootarg.mcc.dr_partner 537403323
```

6. If using encryption with external key manager, set the required bootargs:

```
setenv bootarg.kmip.init.ipaddr
setenv bootarg.kmip.kmip.init.netmask
setenv bootarg.kmip.kmip.init.gateway
setenv bootarg.kmip.kmip.init.interface
```

What's next?

[Reassign the root aggregate disks.](#)

Reassign the root aggregate disks to the new MetroCluster IP controller module

Reassign the root aggregate disks to the new controller module, using the system IDs that you gathered earlier.

About this task

This task is performed in Maintenance mode.

The old system IDs were identified in [Gather information before the upgrade](#).

The examples in this procedure use controllers with the following system IDs:

Node	Old system ID	New system ID
node_B_1	4068741254	1574774970

Steps

1. Cable all other connections to the new controller modules (FC-VI, storage, cluster interconnect, and so on).
2. Halt the system and boot to Maintenance mode from the LOADER prompt:

```
boot_ontap maint
```

3. Display the disks owned by node_B_1-old:

```
disk show -a
```

The command output shows the system ID of the new controller module (1574774970). However, the root aggregate disks are still owned by the old system ID (4068741254). This example doesn't show drives owned by other nodes in the MetroCluster configuration.



Before you proceed with disk reassignment, verify that the pool0 and pool1 disks that belong to the node's root aggregate are displayed in the `disk show` output. In the following example, the output lists the pool0 and pool1 disks owned by node_B_1-old.

```
*> disk show -a
Local System ID: 1574774970

  DISK          OWNER          POOL  SERIAL NUMBER  HOME
DR HOME
-----
.....
...
rr18:9.126L44 node_B_1-old(4068741254) Pool1 PZHYN0MD
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:9.126L49 node_B_1-old(4068741254) Pool1 PPG3J5HA
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L21 node_B_1-old(4068741254) Pool1 PZHTDSZD
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L2  node_B_1-old(4068741254) Pool0 SOM1J2CF
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L3  node_B_1-old(4068741254) Pool0 SOM0CQM5
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:9.126L27 node_B_1-old(4068741254) Pool0 SOM1PSDW
node_B_1-old(4068741254) node_B_1-old(4068741254)
...
```

4. Reassign the root aggregate disks on the drive shelves to the new controller:

```
disk reassign -s <old-sysid> -d <new-sysid>
```



If your MetroCluster IP system is configured with Advanced Disk Partitioning, you must include the DR partner system ID by running the `disk reassign -s old-sysid -d new-sysid -r dr-partner-sysid` command.

The following example shows reassignment of drives:

```

*> disk reassign -s 4068741254 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? Jul 14 19:23:49
[localhost:config.bridge.extra.port:error]: Both FC ports of FC-to-SAS
bridge rtp-fc02-41-rr18:9.126L0 S/N [FB7500N107692] are attached to this
controller.
y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 4068741254.
Do you want to continue (y/n)? y

```

5. Check that all disks are reassigned as expected:

```
disk show
```

```

*> disk show
Local System ID: 1574774970

  DISK          OWNER                                POOL  SERIAL NUMBER  HOME
DR HOME
-----
rr18:8.126L18  node_B_1-new(1574774970)  Pool1  PZHYN0MD
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L49  node_B_1-new(1574774970)  Pool1  PPG3J5HA
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L21  node_B_1-new(1574774970)  Pool1  PZHTDSZD
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L2   node_B_1-new(1574774970)  Pool0  SOM1J2CF
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L29  node_B_1-new(1574774970)  Pool0  SOM0CQM5
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L1   node_B_1-new(1574774970)  Pool0  SOM1PSDW
node_B_1-new(1574774970)  node_B_1-new(1574774970)
*>

```

6. Display the aggregate status:

```
aggr status
```

```
*> aggr status
      Aggr                State      Status      Options
aggr0_node_b_1-root      online    raid_dp, aggr  root, nosnap=on,
                        mirrored
mirror_resync_priority=high(fixed)
                        fast zeroed
                        64-bit
```

7. Repeat the above steps on the partner node (node_B_2-new).

What's next?

[Boot the new controllers and restore LIF configuration.](#)

Boot the new MetroCluster IP controllers and restore LIF configuration

Boot the new controllers and verify that LIFs are hosted on appropriate nodes and ports before resuming the operation by using the `system controller replace resume` command.

Boot the new controllers

Boot the new controllers, verify that the bootarg variables are correct and if needed, perform the encryption recovery steps.

About this task

This task must be performed on all the new controllers.

Steps

1. Halt the node:

```
halt
```

2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr <ip-address>
```

```
setenv bootarg.kmip.init.netmask <netmask>
```

```
setenv bootarg.kmip.init.gateway <gateway-address>
```

```
setenv bootarg.kmip.init.interface <interface-id>
```

3. Display the boot menu:

```
boot_ontap menu
```

4. If root encryption is used, select the boot menu option for your key management configuration.

If you are using...	Select this boot menu option...
Onboard key management	Option "10" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.
External key management	Option "11" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.

5. From the boot menu, run option "6".



Option "6" reboots the node twice before the process completes.

Respond with "y" to the system ID change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...  
  
Rebooting to load the restored env file...
```

During one of the reboots after option "6", the confirmation prompt `Override system ID? {y|n}` appears. Enter `y`.

6. If root encryption is used, select the boot menu option again for your key management configuration.

If you are using...	Select this boot menu option...
Onboard key management	Option "10" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.
External key management	Option "11" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.

Depending on the key manager setting, perform the recovery procedure by selecting option "10" or option "11", followed by option "6" at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option "1" (normal boot).

7. Boot the nodes:

```
boot_ontap
```

8. Wait for the replaced nodes to boot up.

If either node is in takeover mode, perform a giveback using the `storage failover giveback` command.

9. Verify that all ports are in a broadcast domain:

a. View the broadcast domains:

```
network port broadcast-domain show
```

b. If a new broadcast domain is created for the data ports on the newly upgraded controllers, delete the broadcast domain:



Only delete the new broadcast domain. Do not delete any of the broadcast domains that existed before starting the upgrade.

```
broadcast-domain delete -broadcast-domain <broadcast_domain_name>
```

c. Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

d. Add the physical port that will host the intercluster LIFs to the corresponding broadcast domain.

e. Modify intercluster LIFs to use the new physical port as home port.

f. After the intercluster LIFs are up, check the cluster peer status and re-establish cluster peering as needed.

You might need to reconfigure cluster peering.

[Create a cluster peer relationship](#)

g. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Create a VLAN](#)

[Combine physical ports to create interface groups](#)

h. Verify that the partner cluster is reachable and that the configuration is successfully resynchronized on the partner cluster:

```
metrocluster switchback -simulate true
```

10. If encryption is used, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> <p>For more information, see Restore onboard key management encryption keys.</p>

External key management

```
security key-manager external restore
-vserver <svm-name> -node <node-name>
-key-server
<host_name|IP_address:port> -key-id
<key_id> -key-tag key_tag <node-name>
```

For more information, see [Restore external key management encryption keys](#).

11. Verify that the MetroCluster is configured correctly. Check the node status:

```
metrocluster node show
```

Verify that the new nodes (site_B) are in **Waiting for switchback state** from site_A.

Verify and restore LIF configuration

Verify that LIFs are hosted on appropriate nodes before you proceed with the automated switchback operation.

About this task

- This task is performed on site_B.



You must verify that the location of the data LIFs is correct on the new nodes before you perform a switchback. When you switchback the configuration, ONTAP attempts to resume traffic on the home port used by the LIFs. I/O failure can occur when the home port connection to the switch port and VLAN is incorrect.

Steps

1. Verify that LIFs are hosted on the appropriate node and ports before switchback.

a. Change to the advanced privilege level:

```
set -privilege advanced
```

b. Display the LIFs, and confirm that each data LIF is using the correct home port:

```
network interface show
```

c. Modify any LIFs that aren't using the correct home port:

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-port
<port-id>
```

If the command returns an error, you can override the port configuration:

```
vserver config override -command "network interface modify -vserver <svm-
name> -home-port <active_port_after_upgrade> -lif <lif_name> -home-node
<new_node_name>"
```

When entering the network interface modify command within the `vserver config override` command, you cannot use the tab autocomplete feature. You can create the network interface modify using autocomplete and then enclose it in the `vserver config override` command.

d. Confirm that all data LIFs are now on the correct home port:

```
network interface show
```

e. Return to the admin privilege level:

```
set -privilege admin
```

2. Revert the interfaces to their home node:

```
network interface revert * -vserver <svm-name>
```

Perform this step on all SVMs as required.

3. Resume the operation:

```
system controller replace resume
```

What's next?

[Complete the controller upgrade.](#)

Complete the MetroCluster IP controller upgrade

Complete the automated controller upgrade process by verifying the network reachability and restoring any monitoring configuration.

Verify network reachability

The automation operation runs verification system checks and then pauses so you can verify the network reachability. After verification, the resource regain phase is initiated and the automation operation executes switchback at site A and pauses at the post upgrade checks. After you resume the automation operation, it performs the post upgrade checks and if no errors are detected, marks the upgrade as complete.

Steps

1. Verify the network reachability by following the console message.

2. After you complete the verification, resume the operation:

```
system controller replace resume
```

3. The automation operation performs `heal-aggregate`, `heal-root-aggregate`, and `switchback` operations at site A, and the post upgrade checks. When the operation pauses, manually check the SAN LIF status and verify the network configuration by following the console message.

4. After you complete the verification, resume the operation:

```
system controller replace resume
```

5. Check the post upgrade checks status:

```
system controller replace show
```

If the post upgrade checks didn't report any errors, the upgrade is complete.

6. After you complete the controller upgrade, log in at site B and verify that the replaced controllers are configured correctly.

Upgrade the nodes on cluster_A

You must repeat the upgrade tasks to upgrade the nodes on cluster_A at site A.

Steps

1. Repeat the steps to upgrade the nodes on cluster_A, beginning with [Prepare for the upgrade](#).

When you repeat the procedure, all example references to the clusters and nodes are reversed.

Re-add internal drives to the new controller

If you upgraded from a system that only has external drives to a system that has external and internal drives (disks and controllers in the same chassis), you can add or re-seat the disks that you removed or unseated from the internal slots of the new system. You can do this at any time after the upgrade is completed on both sites and the cluster is in a healthy state.

After you re-add or re-seat the drives, they can be used in ONTAP as required.



This task only applies to certain upgrade combinations. Refer to [remove internal drives from the chassis on the new controller](#) for more information.

Reconfigure ONTAP Mediator

Manually configure ONTAP Mediator which was automatically removed before you started the upgrade.

1. Use the steps in [Configure ONTAP Mediator from a MetroCluster IP configuration](#).

Restore Tiebreaker monitoring

If the MetroCluster configuration was previously configured for monitoring by the Tiebreaker software, you can restore the Tiebreaker connection.

1. Use the steps in [Add MetroCluster configurations](#).

Configure end-to-end encryption

If it is supported on your system, you can encrypt back-end traffic, such as NVlog and storage replication data, between the MetroCluster IP sites. Refer to [Configure end-to-end encryption](#) for more information.

Upgrade controllers in MetroCluster IP using switchover and switchback (ONTAP 9.8 and later)

Workflow for MetroCluster IP controller upgrades using switchover and switchback (ONTAP 9.8 and later)

Beginning with ONTAP 9.8, you can use the MetroCluster switchover operation to provide nondisruptive service to clients while the controller modules on the partner cluster are upgraded. Other components (such as storage shelves or switches) cannot be upgraded

as part of this procedure.

About this workflow

You can use this workflow to upgrade MetroCluster IP controllers using switchover and switchback on systems running ONTAP 9.8 or later.

1

Prepare to upgrade

Review the supported upgrade combinations and requirements, and complete the required tasks to prepare your system for the controller upgrade.

2

Upgrade your controllers

Switch over the MetroCluster configuration in order to remove the configuration from the old controllers, rack and install the new controllers, reassign the root aggregate disks, and boot the new controllers before you perform a switchback.

3

Complete the upgrade

Complete the controller upgrade by repeating the upgrade tasks at the second site and restoring any monitoring configuration.

Prepare to upgrade

Requirements for using this MetroCluster IP upgrade procedure

Verify that your system meets all the requirements before performing the controller upgrade.



You must perform this procedure as written each time you upgrade your controllers.

When new platforms are released, there might be new or modified steps that must be followed. For example, when upgrading to platforms introduced in ONTAP 9.15.1 or later, you must [set the required bootarg](#) and perform other additional steps for a successful upgrade.

Platforms supported by this procedure

- The platforms must be running ONTAP 9.8 or later.
- The target (new) platform must be a different model than the original platform.
- You can only upgrade specific platform models using this procedure in a MetroCluster IP configuration.
 - For information on what platform upgrade combinations are supported, review the MetroCluster IP upgrade table in [Choose a controller upgrade procedure](#).

Refer to [Choosing an upgrade or refresh method](#) for additional procedures.

Requirements

- This procedure applies to controller modules in a MetroCluster IP configuration.
- All controllers in the configuration should be upgraded during the same maintenance period.

Operating the MetroCluster configuration with different controller types is not supported outside of this maintenance activity.

- The MetroCluster IP systems must be running the same ONTAP version at both sites.
- The MetroCluster IP switches (switch type, vendor, and model) and firmware version must be supported on the existing and new controllers in your upgrade configuration.

Refer to the [Hardware Universe](#) or the [IMT](#) for supported switches and firmware versions.

- When you upgrade from systems that have more slots or ports than the new system, you need to verify that there are sufficient slots and ports on the new system.

Before you start the upgrade, refer to the [Hardware Universe](#) to verify the number of slots and ports on the new system.

- If it's enabled on your system, [disable end-to-end encryption](#) before performing the upgrade.
- If the new platform has fewer slots than the original system, or if it has fewer or different types of ports, you might need to add an adapter to the new system.
- You reuse the IP addresses, netmasks, and gateways of the original platforms on the new platforms.

The following example names are used in this procedure:

- cluster_A at site_A
 - Before upgrade:
 - node_A_1-old
 - node_A_2-old
 - After upgrade:
 - node_A_1-new
 - node_A_2-new
- cluster_B at site_B
 - Before upgrade:
 - node_B_1-old
 - node_B_2-old
 - After upgrade:
 - node_B_1-new
 - node_B_2-new

What's next?

[Enable console logging.](#)

Enable console logging before the MetroCluster IP controller upgrade

Enable console logging on your devices before performing the controller upgrade.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

What's next?

Review the information in [Set the required bootarg \(for upgrades to systems introduced in 9.15.1 or later\)](#) to confirm whether you need to set a required bootarg on the existing system.

Set the required bootarg (for MetroCluster IP upgrades to systems introduced in ONTAP 9.15.1 or later)

If you are upgrading to a system introduced in ONTAP 9.15.1 or later, you usually need to set a bootarg on the old controllers before you can start the upgrade.



If your upgrade combination is affected, setting a bootarg on all of the old controllers is required for a successful upgrade. Review the information in this section carefully to check if your upgrade combination requires that you set a bootarg, and how to set the correct bootarg for your combination.

Step 1: Determine whether you need to set a bootarg on the old controllers

Before you start the upgrade, determine if you need to set a bootarg on the old controllers using the following information:

- You **must** set a bootarg on the old controllers for supported upgrades to the following systems unless stated otherwise:
 - AFF A70, AFF A90, AFF A1K
 - FAS70, FAS90
 - AFF C80
 - AFF A50, AFF A20, AFF A30
 - AFF C30, AFF C60
 - FAS50
- You do **not** have to set any bootarg on the old controllers if your upgrade is one of the following combinations:
 - From an AFF A70 system to an AFF A90 system
 - From a FAS70 system to a FAS90 system



If your upgrade doesn't require you to set any bootarg, you can skip this task and go directly to [Prepare the system for upgrade](#).

Step 2: Determine the bootarg you need to set on the old controllers

Most affected upgrades require you to set the `hw.cxgbe.toe_keepalive_disable` bootarg on the old controllers. However, certain upgrade paths require you to set the `bootarg.siw.interop_enabled` bootarg instead.

Use the following table to determine which bootarg you need to set for your specific upgrade combination.

For this upgrade...	Set the bootarg...
From AFF A250 to AFF A30	<code>bootarg.siw.interop_enabled</code>
From AFF C250 to AFF C30	<code>bootarg.siw.interop_enabled</code>
From AFF A150 to AFF A20	<code>bootarg.siw.interop_enabled</code>
From AFF A220 to AFF A20	<code>bootarg.siw.interop_enabled</code>
All other upgrades to AFF A70, AFF A90, AFF A1K, FAS70, FAS90, AFF C80, AFF A50, AFF A30, AFF C30, AFF C60, or FAS50 systems	<code>hw.cxgbe.toe_keepalive_disable</code>

Note: You don't have to set any bootarg if your upgrade is from an AFF A70 to AFF A90 system or from a FAS70 to FAS90 system.

Step 3: Set the required bootarg on the old controllers

After you've determined the bootarg required for your upgrade combination, follow the steps to set the bootarg on the old controllers.



You must set the bootarg on all of the old nodes at both sites before you start the upgrade.

Steps

1. Halt one node at both sites and allow its HA partner to perform a storage takeover of the node:

```
halt -node <node_name>
```

2. Set the required bootarg for your upgrade combination. You already determined the bootarg that you need to set by using the table in [determine which bootarg you need to set](#).

hw.cxgbe.toe_keepalive_disable

- a. At the LOADER prompt of the halted node, enter the following:

```
setenv hw.cxgbe.toe_keepalive_disable 1  
  
saveenv  
  
printenv hw.cxgbe.toe_keepalive_disable
```

bootarg.siw.interop_enabled

- a. At the LOADER prompt of the halted node, enter the following:

```
setenv bootarg.siw.interop_enabled 1  
  
saveenv  
  
printenv bootarg.siw.interop_enabled
```

3. Boot the node:

```
boot_ontap
```

4. When the node boots, perform a giveback for the node at the prompt:

```
storage failover giveback -ofnode <node_name>
```

5. Repeat the steps on all nodes in the DR group or DR groups that you are upgrading.

What's next?

[Prepare the system for upgrade.](#)

Prepare the MetroCluster IP system for upgrade

Before making any changes to the existing MetroCluster configuration, check the health of the configuration, prepare the new platforms, and perform other miscellaneous tasks.

Update the MetroCluster switch RCFs before upgrading controllers

Depending on the old and new platform models, you might need to update the MetroCluster switch reference configuration files (RCFs) before you upgrade controllers.

About this task

Perform this task in the following circumstances:

- The switch RCF configuration is not on the minimum version.
- You need to change VLAN IDs used by the back-end MetroCluster connections.

Before you begin

Determine whether you need to update the RCFs before you upgrade your controllers:

- If the switch configuration wasn't configured with the minimum supported RCF version, you need to update the RCFs before you upgrade your controllers:

Switch model	Required RCF version
Cisco 3132Q-V	1.7 or later
Cisco 3232C	1.7 or later
Broadcom BES-53248	1.3 or later
NVIDIA SN2100	2.0 or later

- If both of your old and new platform models are in the following list, you do **not** need to update the VLAN ID before you upgrade controllers:
 - FAS8200 or AFF A300
 - AFF A320
 - FAS9000 or AFF A700
 - AFF A800, AFF C800, ASA A800, or ASA C800

If either of your old or new platform models are not listed above, you must confirm that the MetroCluster interfaces are using a supported VLAN ID. Supported VLAN IDs for the MetroCluster interfaces are: 10, 20, or in the range of 101 to 4096.



- If the VLAN ID is not 10, 20, or in the range of 101 to 4096, you must upgrade the switch RCF before you upgrade controllers.
- The local cluster connections can use any VLAN, they don't need to be in the given range.
- The new RCF that you are upgrading to must use the VLANs 10, 20, or be in the range 101 to 4096. Don't change the VLAN for the local cluster unless it is required.

Steps

1. Prepare the IP switches for the application of the new RCFs.

Follow the steps in the section for your switch vendor:



You should update the switches in the following order: switch_A_1, switch_B_1, switch_A_2, switch_B_2.

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

2. Download and install the RCFs.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom RCFs](#)
- [Download and install the Cisco IP RCFs](#)
- [Download and install the NVIDIA IP RCFs](#)

Map ports from the old nodes to the new nodes

You must verify that the physical ports on node_A_1-old map correctly to the physical ports on node_A_1-new. This allows node_A_1-new to communicate with other nodes in the cluster and with the network after the upgrade.

About this task

When the new node first boots during the upgrade process, it replays the most recent configuration of the old node it's replacing. When you boot node_A_1-new, ONTAP attempts to host LIFs on the same ports that were used on node_A_1-old. This means that you have to adjust the port and LIF configuration as part of the upgrade so it's compatible with the configuration of the old node. During the upgrade procedure, you perform steps on both the old and new nodes to ensure correct configuration for the cluster, management, and data LIFs

The following table shows examples of configuration changes related to the port requirements of the new nodes.

Cluster interconnect physical ports		
Old controller	New controller	Required action
e0a, e0b	e3a, e3b	No matching port. After the upgrade, you must recreate the cluster ports.
e0c, e0d	e0a,e0b,e0c,e0d	e0c and e0d are matching ports. You don't have to change the configuration, but after the upgrade you can spread your cluster LIFs across the available cluster ports.

Steps

1. Determine what physical ports are available on the new controllers and what LIFs can be hosted on the ports.

The controller's port usage depends on the platform module and which switches you will use in the MetroCluster IP configuration. You can gather the port usage of the new platforms from the [Hardware Universe](#).

2. Plan your port usage and fill in the following tables for reference for each of the new nodes.

You will refer to the table as you carry out the upgrade procedure.

LIF	node_A_1-old			node_A_1-new		
	Ports	IPspaces	Broadcast domains	Ports	IPspaces	Broadcast domains

Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

Netboot the new controllers

After you install the new nodes, you need to netboot to ensure the new nodes are running the same version of ONTAP as the original nodes. The term netboot means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

Steps

1. Netboot the new controllers:
 - a. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
 - b. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `ontap-version_image.tgz` file on a web-accessible directory.
 - c. Change to the web-accessible directory and verify that the files you need are available.

Your directory listing should contain a netboot folder with a kernel file:

```
_ontap-version_image.tgz
```

You don't need to extract the `_ontap-version_image.tgz` file.

- d. At the `LOADER` prompt, configure the netboot connection for a management LIF:

If IP addressing is...	Then...
DHCP	Configure the automatic connection: <code>ifconfig e0M -auto</code>
Static	Configure the manual connection: <code>ifconfig e0M -addr=<i>ip_addr</i> - mask=<i>netmask</i> -gw=<i>gateway</i></code>

- e. Perform the netboot.

```
netboot http://_web_server_ip/path_to_web-accessible_directory/ontap-  
version_image.tgz
```

- f. From the boot menu, select option **(7) Install new software first** to download and install the new software image to the boot device.

Disregard the following message:

```
"This procedure is not supported for Non-Disruptive Upgrade on an HA pair". It  
applies to nondisruptive upgrades of software, not to upgrades of controllers.
```

- g. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL of the image file:

```
http://web_server_ip/path_to_web-accessible_directory/ontap-  
version_image.tgz
```

- h. Enter the user name and password if applicable, or press `Enter` to continue.

- i. Enter `n` to skip the backup recovery when you see a prompt similar to the following:

```
Do you want to restore the backup configuration now? {y|n} n
```

- j. Reboot by entering `y` when you see a prompt similar to the following:

```
The node must be rebooted to start using the newly installed  
software. Do you want to reboot now? {y|n} y
```



You must reboot the node in order to use the newly installed software.

Clear the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing

configuration.

Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

Verify MetroCluster health before site upgrade

You verify the health and connectivity of the MetroCluster configuration before performing the upgrade.



After you upgrade the controllers at the first site and before you upgrade the second, running `metrocluster check run` followed by `metrocluster check show` returns an error in the `config-replication` field. This error indicates an NVRAM size mismatch between the nodes at each site and it's the expected behavior when there are different platform models on both sites. You can ignore the error until the controller upgrade is completed for all nodes in the DR group.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the nodes are multipathed:

```
node run -node <node_name> sysconfig -a
```

Issue this command for each node in the MetroCluster configuration.

- b. Verify that there are no broken disks in the configuration:

```
storage disk show -broken
```

Issue this command on each node in the MetroCluster configuration.

- c. Check for any health alerts:

```
system health alert show
```

Issue this command on each cluster.

- d. Verify the licenses on the clusters:

```
system license show
```

Issue this command on each cluster.

- e. Verify the devices connected to the nodes:

```
network device-discovery show
```

Issue this command on each cluster.

- f. Verify that the time zone and time is set correctly on both sites:

```
cluster date show
```

Issue this command on each cluster. You can use the `cluster date` commands to configure the time and time zone.

2. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

- a. Confirm the MetroCluster configuration and that the operational mode is `normal`:

```
metrocluster show
```

- b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

- c. Issue the following command:

```
metrocluster check run
```

- d. Display the results of the MetroCluster check:

```
metrocluster check show
```

3. Check the MetroCluster cabling with the Config Advisor tool.

- a. Download and run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- b. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Gather information before the upgrade

Before upgrading, you must gather information for each of the nodes, and, if necessary, adjust the network broadcast domains, remove any VLANs and interface groups, and gather encryption information.

Steps

1. Record the physical cabling for each node, labelling cables as needed to allow correct cabling of the new nodes.
2. Gather interconnect, port, and LIF information for each node.

Gather the output of the following commands for each node:

- `metrocluster interconnect show`
- `metrocluster configuration-settings connection show`
- `network interface show -role cluster,node-mgmt`
- `network port show -node <node_name> -type physical`
- `network port vlan show -node <node_name>`
- `network port ifgrp show -node <node_name> -instance`
- `network port broadcast-domain show`
- `network port reachability show -detail`
- `network ipspace show`
- `volume show`
- `storage aggregate show`
- `system node run -node <node_name> sysconfig -a`
- `aggr show -r`
- `disk show`
- `system node run <node-name> disk show`
- `vol show -fields type`
- `vol show -fields type , space-guarantee`
- `vserver fcp initiator show`
- `storage disk show`
- `metrocluster configuration-settings interface show`

3. Gather the UUIDs for the site_B (the site whose platforms are currently being upgraded):

```
metrocluster node show -fields node-cluster-uuid, node-uuid
```

These values must be configured accurately on the new site_B controller modules to ensure a successful upgrade. Copy the values to a file so that you can copy them into the commands later in the upgrade process.

The following example shows the command output with the UUIDs:

```

cluster_B::> metrocluster node show -fields node-cluster-uuid, node-uuid
(metrocluster node show)
dr-group-id cluster      node      node-uuid
node-cluster-uuid
-----
1              cluster_A node_A_1 f03cb63c-9a7e-11e7-b68b-00a098908039
ee7db9d5-9a82-11e7-b68b-00a098908039
1              cluster_A node_A_2 aa9a7a7a-9a81-11e7-a4e9-00a098908c35
ee7db9d5-9a82-11e7-b68b-00a098908039
1              cluster_B node_B_1 f37b240b-9ac1-11e7-9b42-00a098c9e55d
07958819-9ac6-11e7-9b42-00a098c9e55d
1              cluster_B node_B_2 bf8e3f8f-9ac4-11e7-bd4e-00a098ca379f
07958819-9ac6-11e7-9b42-00a098c9e55d
4 entries were displayed.
cluster_B::~*

```

NetApp recommends that you record the UUIDs in a table similar to the following:

Cluster or node	UUID
cluster_B	07958819-9ac6-11e7-9b42-00a098c9e55d
node_B_1	f37b240b-9ac1-11e7-9b42-00a098c9e55d
node_B_2	bf8e3f8f-9ac4-11e7-bd4e-00a098ca379f
cluster_A	ee7db9d5-9a82-11e7-b68b-00a098908039
node_A_1	f03cb63c-9a7e-11e7-b68b-00a098908039
node_A_2	aa9a7a7a-9a81-11e7-a4e9-00a098908c35

4. If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

Gather the output of the following commands:

- `fcg adapter show -instance`
- `fcg interface show -instance`
- `iscsi interface show`
- `ucadmin show`

5. If the root volume is encrypted, collect and save the passphrase used for the key manager:

```
security key-manager backup show
```

6. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.

For additional information, see [Back up onboard key management information manually](#).

- a. If Onboard Key Manager is configured:

```
security key-manager onboard show-backup
```

You need the passphrase later in the upgrade procedure.

- b. If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
security key-manager key query
```

7. Gather the system IDs of the existing nodes:

```
metrocluster node show -fields node-systemid,ha-partner-systemid,dr-partner-
systemid,dr-auxiliary-systemid
```

The following output shows the reassigned drives.

```
::> metrocluster node show -fields node-systemid,ha-partner-systemid,dr-
partner-systemid,dr-auxiliary-systemid

dr-group-id cluster      node      node-systemid ha-partner-systemid dr-
partner-systemid dr-auxiliary-systemid
-----
-----
1          cluster_A node_A_1  537403324    537403323
537403321          537403322
1          cluster_A node_A_2  537403323    537403324
537403322          537403321
1          cluster_B node_B_1  537403322    537403321
537403323          537403324
1          cluster_B node_B_2  537403321    537403322
537403324          537403323
4 entries were displayed.
```

Remove Mediator or Tiebreaker monitoring

Before the upgrading the platforms, you must remove monitoring if the MetroCluster configuration is monitored with the Tiebreaker or Mediator utility.

Steps

1. Collect the output for the following command:

```
storage iscsi-initiator show
```

2. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

If you are using...	Use this procedure...
Tiebreaker	Removing MetroCluster Configurations
Mediator	Issue the following command from the ONTAP prompt: <pre>metrocluster configuration-settings mediator remove</pre>
Third-party applications	Refer to the product documentation.

Send a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. Log in to the cluster.
2. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

The `maintenance-window-in-hours` parameter specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

3. Repeat these steps on the partner site.

What's next?

[Switch over the MetroCluster configuration.](#)

Upgrade your controllers

Switch over the MetroCluster IP configuration

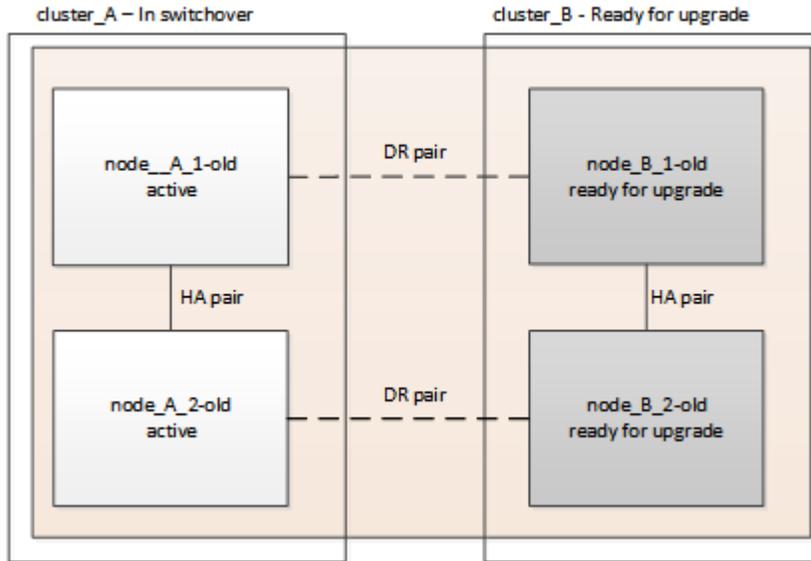
You switch over the configuration to site_A so that the platforms on site_B can be upgraded.

About this task

This task must be performed on site_A.

After you complete this task:

- cluster_A is active and serving data for both sites.
- cluster_B is inactive and ready to begin the upgrade process.



Steps

1. Switch over the MetroCluster configuration to site_A so that site_B's nodes can be upgraded:

a. Issue the following command on cluster_A:

```
metrocluster switchover -controller-replacement true
```

The operation can take several minutes to complete.

b. Monitor the switchover operation:

```
metrocluster operation show
```

c. After the operation is complete, confirm that the nodes are in switchover state:

```
metrocluster show
```

d. Check the status of the MetroCluster nodes:

```
metrocluster node show
```

Automatic healing of aggregates after negotiated switchover is disabled during a controller upgrade.

What's next?

[Remove interface configurations and uninstall the old controllers.](#)

Remove interface configurations and uninstall the old MetroCluster IP controllers

Verify the correct LIF placement. Then remove the VLANs and interface groups on the

old controllers and physically uninstall the controllers.

About this task

- You perform these steps on the old controllers (node_B_1-old, node_B_2-old).
- You require the information you gathered in [Map ports from the old nodes to the new nodes](#) for use in this procedure.

Steps

1. Boot the old nodes and log in to the nodes:

```
boot_ontap
```

2. If the system you are upgrading to uses **shared cluster/HA ports**, verify that the MetroCluster IP interfaces are using supported IP addresses.

Use the following information to determine whether the new system uses shared cluster/HA ports:

Shared cluster/HA ports

The systems listed in the following table use shared cluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A20• AFF A30• AFF C30• AFF A50• AFF C60• AFF C80• AFF A70• AFF A90• AFF A1K	<ul style="list-style-type: none">• FAS50• FAS70• FAS90

Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A150, ASA A150• AFF A220• AFF C250, ASA C250• AFF A250, ASA A250• AFF A300• AFF A320• AFF C400, ASA C400• AFF A400, ASA A400• AFF A700• AFF C800, ASA C800• AFF A800, ASA A800• AFF A900, ASA A900	<ul style="list-style-type: none">• FAS2750• FAS500f• FAS8200• FAS8300• FAS8700• FAS9000• FAS9500

a. Verify the IP addresses of the MetroCluster interfaces on the old controllers:

```
metrocluster configuration-settings interface show
```

b. If the MetroCluster interfaces are using 169.254.17.x or 169.254.18.x IP addresses, refer to [the Knowledge Base article "How to modify the properties of a MetroCluster IP interface"](#) to modify the interface IP addresses before you proceed with the upgrade.



Upgrading to any system that uses **shared cluster/HA ports** isn't supported if the MetroCluster interfaces are configured with 169.254.17.x or 169.254.18.x IP addresses.

3. Modify the intercluster LIFs on the old controllers to use a different home port than the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers.



This step is required for a successful upgrade.

The intercluster LIFs on the old controllers must use a different home port than the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers. For example, when you upgrade to AFF A90 controllers, the HA interconnect ports are e1a and e7a, and the MetroCluster IP DR interconnect ports are e2b and e3b. You must move the intercluster LIFs on the old controllers if they are hosted on ports e1a, e7a, e2b, or e3b.

For port distribution and allocation on the new nodes, refer to the [Hardware Universe](#).

- a. On the old controllers, view the intercluster LIFs:

```
network interface show -role intercluster
```

Take one of the following actions depending on whether the intercluster LIFs on the old controllers use the same ports as the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers.

If the intercluster LIFs...	Go to...
Use the same home port	Substep b
Use a different home port	Step 4

- b. Modify the intercluster LIFs to use a different home port:

```
network interface modify -vserver <vserver> -lif <intercluster_lif> -home
-port <port-not-used-for-ha-interconnect-or-mcc-ip-dr-interconnect-on-new-
nodes>
```

- c. Verify that all intercluster LIFs are on their new home ports:

```
network interface show -role intercluster -is-home false
```

The command output should be empty, indicating that all intercluster LIFs are on their respective home ports.

- d. Revert any LIFs that aren't on their home ports:

```
network interface revert -lif <intercluster_lif>
```

Repeat the command for each intercluster LIF that isn't on the home port.

4. Assign the home port of all data LIFs on the old controller to a common port that is the same on both the old and new controller modules.



If the old and new controllers don't have a common port, you don't need to modify the data LIFs. Skip this step and go directly to [Step 5](#).

a. Display the LIFs:

```
network interface show
```

All data LIFs including SAN and NAS are admin up and operationally down because those are up at switchover site (cluster_A).

b. Review the output to find a common physical network port that is the same on both the old and new controllers that is not used as a cluster port.

For example, e0d is a physical port on old controllers and is also present on new controllers. e0d is not used as a cluster port or otherwise on the new controllers.

For port usage for platform models, see the [Hardware Universe](#)

c. Modify all data LIFS to use the common port as the home port:

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-port <port-id>
```

In the following example, this is "e0d".

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

5. Modify broadcast domains to remove the VLAN and physical ports that need to be deleted:

```
broadcast-domain remove-ports -broadcast-domain <broadcast-domain-name> -ports <node-name:port-id>
```

Repeat this step for all VLAN and physical ports.

6. Remove any VLAN ports using cluster ports as member ports and interface groups using cluster ports as member ports.

a. Delete VLAN ports:

```
network port vlan delete -node <node_name> -vlan-name <portid-vlandid>
```

For example:

```
network port vlan delete -node node1 -vlan-name e1c-80
```

b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node <node_name> -ifgrp <interface-group-name> -port <portid>
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp ala -port e0d
```

c. Remove VLAN and interface group ports from the broadcast domain:

```
network port broadcast-domain remove-ports -ip-space <ip-space> -broadcast
-domain <broadcast-domain-name> -ports
<nodename:portname,nodename:portname>,..
```

d. Modify interface group ports to use other physical ports as member, as needed:

```
ifgrp add-port -node <node_name> -ifgrp <interface-group-name> -port <port-
id>
```

7. Halt the nodes to the `LOADER` prompt:

```
halt -inhibit-takeover true
```

8. Connect to the serial console of the old controllers (`node_B_1-old` and `node_B_2-old`) at `site_B` and verify it is displaying the `LOADER` prompt.

9. Gather the bootarg values:

```
printenv
```

10. Disconnect the storage and network connections on `node_B_1-old` and `node_B_2-old`. Label the cables so that you can reconnect them to the new nodes.

11. Disconnect the power cables from `node_B_1-old` and `node_B_2-old`.

12. Remove the `node_B_1-old` and `node_B_2-old` controllers from the rack.

What's next?

[Set up the new controllers.](#)

Set up the new MetroCluster IP controllers

Rack and cable the new MetroCluster IP controllers.

Steps

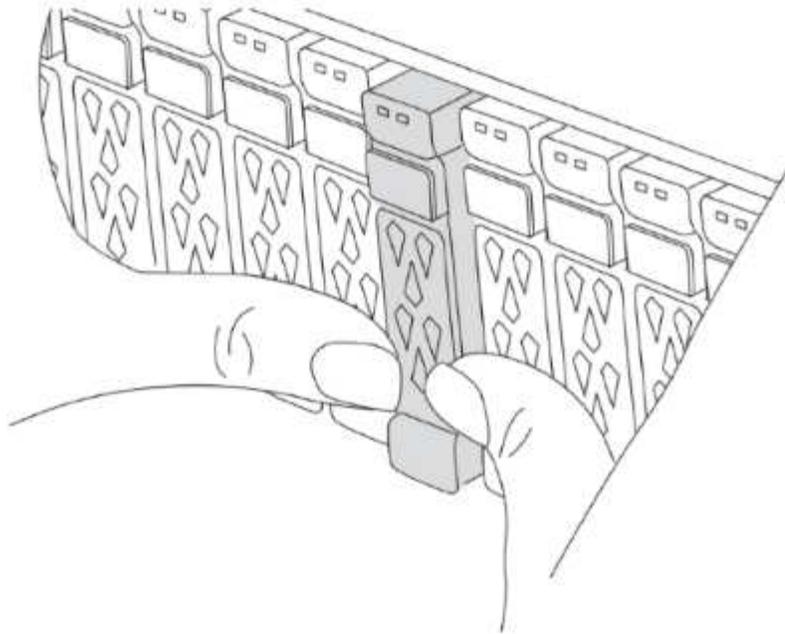
1. Plan out the positioning of the new controller modules and storage shelves as needed.

The rack space depends on the platform model of the controller modules, the switch types, and the number of storage shelves in your configuration.

2. Properly ground yourself.

3. If your upgrade requires replacement of the controller modules, for example, upgrading from an AFF A800 to an AFF A90 system or from an AFF C800 to an AFF C80 system, you must remove the controller module from the chassis when you replace the controller module. For all other upgrades, skip to [Step 4](#).

On the front of the chassis, use your thumbs to firmly push each drive in until you feel a positive stop. This confirms that the drives are firmly seated against the chassis midplane.



4. Install the controller modules.

The installation steps you follow depend on whether your upgrade requires replacement of the controller modules, or if IOM modules are required to convert the old controllers to an external shelf.

If you are upgrading...	Follow the steps for ...
<ul style="list-style-type: none"> • An AFF A150 to an AFF A20 system • An AFF A220 to an AFF A20 system 	Controller to external shelf conversion
<ul style="list-style-type: none"> • An AFF A800 to an AFF A90 system • An AFF C800 to an AFF C80 system • An AFF A250 to an AFF A30 system • An AFF C250 to an AFF C30 system • An AFF A70 to an AFF A90 system 	Controller module replacement
Any other controller upgrade combinations	All other upgrades

Controller to external shelf conversion

If your original MetroCluster IP controllers are AFF A150 or AFF A220 models, you can convert the AFF A150 or AFF A220 HA pair to a DS224C drive shelf and then attach it to the new nodes.

For example, when upgrading from an AFF A150 or AFF A220 system to an AFF A20 system, you can convert the AFF A150 or AFF A220 HA pair to a DS224C shelf by swapping the AFF A150 or AFF A220 controller modules with IOM12 modules.

Steps

1. Replace the controller modules in the node you are converting with IOM12 shelf modules.

[Hardware Universe](#)

2. Set the drive shelf ID.

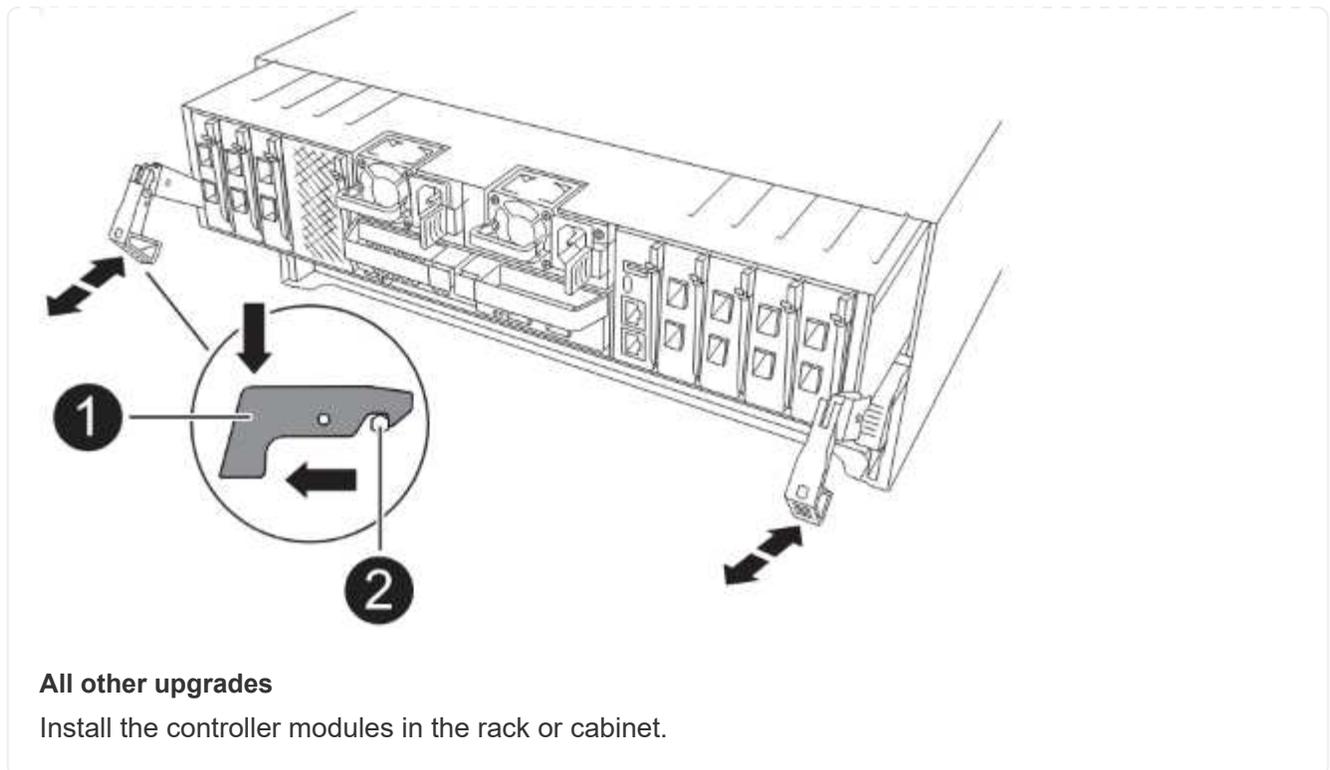
Each drive shelf, including the chassis, requires a unique ID.

3. Reset other drive shelf IDs as needed.
4. Power off the shelves.
5. Cable the converted drive shelf to a SAS port on the new system, and, if you are using out-of-band ACP cabling, to the ACP port on the new node.
6. Turn on the power to the converted drive shelf and any other drive shelves attached to the new nodes.
7. Turn on the power to the new nodes, and then interrupt the boot process on each node by pressing Ctrl-C to access the boot environment prompt.

Controller module replacement

Installing the new controllers separately is not applicable for upgrades of integrated systems with disks and controllers in the same chassis, for example, from an AFF A800 system to an AFF A90 system. You must swap the new controller modules and I/O cards after powering off the old controllers, as shown in the image below.

The following example image is for representation only, the controller modules and I/O cards can vary between systems.



5. Cable the controllers' power, serial console, and management connections as described in [Cable the MetroCluster IP switches](#).

Don't connect any other cables that were disconnected from old controllers at this time.

[ONTAP Hardware Systems Documentation](#)

6. Power up the new nodes and boot them to Maintenance mode.

What's next?

[Restore the HBA configuration and set the HA state.](#)

Restore the HBA configuration and set the HA state of the MetroCluster IP controller and chassis

Configure the HBA cards in the controller module and verify and set the HA state of the controller and chassis.

Restore the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site.

Steps

1. In Maintenance mode, configure the settings for any HBAs in the system:
 - a. Check the current settings of the ports: `ucadmin show`
 - b. Update the port settings as needed.

If you have this type of HBA and desired mode...	Use this command...
--	---------------------

CNA FC	<code>ucadmin modify -m fc -t initiator <adapter-name></code>
CNA Ethernet	<code>ucadmin modify -mode cna <adapter-name></code>
FC target	<code>fcadmin config -t target <adapter-name></code>
FC initiator	<code>fcadmin config -t initiator <adapter-name></code>

2. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the `LOADER` prompt.

3. Boot the node back into Maintenance mode to apply the configuration changes:

```
boot_ontap maint
```

4. Verify the changes:

If you have this type of HBA...	Use this command...
CNA	<code>ucadmin show</code>
FC	<code>fcadmin show</code>

Set the HA state on the new controllers and chassis

You must verify the HA state of the controllers and chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be `mccip`.

2. If the displayed system state of the controller or chassis isn't correct, set the HA state:

```
ha-config modify controller mccip
```

```
ha-config modify chassis mccip
```

3. Verify and modify the Ethernet ports connected to NS224 shelves or storage switches.

a. Verify the Ethernet ports connected to NS224 shelves or storage switches:

```
storage port show
```

b. Set all Ethernet ports connected to Ethernet shelves or storage switches, including shared switches for storage and cluster, to `storage` mode:

```
storage port modify -p <port> -m storage
```

Example:

```
*> storage port modify -p e5b -m storage
Changing NVMe-oF port e5b to storage mode
```



This must be set on all affected ports for a successful upgrade.

Disks from the shelves attached to the Ethernet ports are reported in the `sysconfig -v` output.

Refer to the [Hardware Universe](#) for information on the storage ports for the system you are upgrading to.

c. Verify that `storage` mode is set and confirm that the ports are in the online state:

```
storage port show
```

4. Halt the node: `halt`

The node should stop at the `LOADER>` prompt.

5. On each node, check the system date, time, and time zone: `show date`

6. If necessary, set the date in UTC or GMT: `set date <mm/dd/yyyy>`

7. Check the time by using the following command at the boot environment prompt: `show time`

8. If necessary, set the time in UTC or GMT: `set time <hh:mm:ss>`

9. Save the settings: `saveenv`

10. Gather environment variables: `printenv`

Remove internal drives from the chassis on the new controller

When you upgrade from a system that only has external drives to a system that has external and internal drives (disks and controllers in the same chassis), you need to remove or unseat all internal drives from the new system until you have completed the upgrade.



This task is mandatory for a successful controller upgrade on affected systems.

To determine if your upgrade combination is affected, refer to the table in [Supported controller upgrades](#). If your upgrade combination is marked with **Note 3**, you must remove or unseat the internal drives from the new

system.

After you complete this task, no internal drives should be accessible. You'll add the drives to the new controller later in the procedure.

What's next?

[Update the switch RCFs and set the MetroCluster IP bootarg values.](#)

Update the switch RCFs and set the MetroCluster IP bootarg values

Update the switch reference configuration files (RCFs) for the new platforms and set the MetroCluster IP bootarg values on the controller modules.

Update the switch RCFs to accommodate the new platforms

You must update the switches to a configuration that supports the new platform models.

About this task

You perform this task at the site containing the controllers that are currently being upgraded. In the examples shown in this procedure we are upgrading site_B first.

The switches at site_A will be upgraded when the controllers on site_A are upgraded.

Steps

1. Prepare the IP switches for the application of the new RCFs.

Follow the steps in the section for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

2. Download and install the RCFs.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom RCFs](#)
- [Download and install the Cisco IP RCFs](#)
- [Download and install the NVIDIA IP RCFs](#)

Set the MetroCluster IP bootarg variables

You must configure certain MetroCluster IP bootarg values on the new controller modules. The bootarg values must match those configured on the old controller modules.

About this task

- You use the UUIDs and system IDs identified earlier in the upgrade procedure in [Gather information before the upgrade](#).
- Depending on your platform model, you can specify the VLAN ID using the `-vlan-id` parameter. The following platforms do not support the `-vlan-id` parameter:
 - FAS8200 and AFF A300

- AFF A320
- FAS9000 and AFF A700
- AFF C800, ASA C800, AFF A800, and ASA A800

All other platforms support the `-vlan-id` parameter.

- The MetroCluster bootarg values you set depend on whether your new system uses shared cluster/HA ports or shared MetroCluster/HA ports.

Shared cluster/HA ports

The systems listed in the following table use shared cluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A20• AFF A30• AFF C30• AFF A50• AFF C60• AFF C80• AFF A70• AFF A90• AFF A1K	<ul style="list-style-type: none">• FAS50• FAS70• FAS90

Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A150, ASA A150• AFF A220• AFF C250, ASA C250• AFF A250, ASA A250• AFF A300• AFF A320• AFF C400, ASA C400• AFF A400, ASA A400• AFF A700• AFF C800, ASA C800• AFF A800, ASA A800• AFF A900, ASA A900	<ul style="list-style-type: none">• FAS2750• FAS500f• FAS8200• FAS8300• FAS8700• FAS9000• FAS9500

Steps

1. At the `LOADER>` prompt, set the following bootargs on the new nodes at `site_B`:

The steps you follow depend on the ports used by the new platform model.

Systems that use shared cluster/HA ports

- a. Set the following bootargs:

```
setenv bootarg.mcc.port_a_ip_config <local-IP-address/local-IP-  
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id>
```

```
setenv bootarg.mcc.port_b_ip_config <local-IP-address/local-IP-  
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id>
```



If the interfaces are using a default VLAN ID, the `vlan-id` parameter is not required.

The following example sets the values for node_B_1-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12,130
```

The following example sets the values for node_B_2-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,0,172.17.26.12,172.17.26.13,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,0,172.17.27.12,172.17.27.13,130
```

The following example sets the values for node_B_1-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12
```

The following example sets the values for node_B_2-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,0,172.17.26.12,172.17.26.13  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,0,172.17.27.12,172.17.27.13
```

Systems that use shared MetroCluster/HA ports

- a. Set the following bootargs:

```
setenv bootarg.mcc.port_a_ip_config <local-IP-address/local-IP-  
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-  
address,vlan-id>
```

```
setenv bootarg.mcc.port_b_ip_config <local-IP-address/local-IP-  
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-  
address,vlan-id>
```



If the interfaces are using a default VLAN ID, the `vlan-id` parameter is not required.

The following example sets the values for node_B_1-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12,130
```

The following example sets the values for node_B_2-new using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13,130
```

The following example sets the values for node_B_1-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12
```

The following example sets the values for node_B_2-new using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13
setenv bootarg.mcc.port_b_ip_config
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13
```

2. At the new nodes' LOADER prompt, set the UUIDs:

```
setenv bootarg.mgwd.partner_cluster_uuid <partner-cluster-UUID>
```

```
setenv bootarg.mgwd.cluster_uuid <local-cluster-UUID>
```

```
setenv bootarg.mcc.pri_partner_uuid <DR-partner-node-UUID>
```

```
setenv bootarg.mcc.aux_partner_uuid <DR-aux-partner-node-UUID>
```

```
setenv bootarg.mcc.iscsi.node_uuid <local-node-UUID>
```

a. Set the UUIDs on node_B_1-new:

The following example shows the commands for setting the UUIDs on node_B_1-new:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.pri_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.aux_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f
setenv bootarg.mcc.iscsi.node_uuid f03cb63c-9a7e-11e7-b68b-
00a098908039
```

b. Set the UUIDs on node_B_2-new:

The following example shows the commands for setting the UUIDs on node_B_2-new:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.pri_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f
setenv bootarg.mcc.aux_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.iscsi.node_uuid aa9a7a7a-9a81-11e7-a4e9-
00a098908c35
```

- Determine whether the original systems were configured for Advanced Drive Partitioning (ADP) by running the following command from the site that is up:

```
disk show
```

The "container type" column displays "shared" in the `disk show` output if ADP is configured. If "container type" has any other value, ADP is not configured on the system. The following example output shows a system configured with ADP:

```
::> disk show
```

Disk Owner	Usable Size	Shelf	Bay	Disk Type	Container Type	Container Name
Info: This cluster has partitioned disks. To get a complete list of spare disk capacity use "storage aggregate show-spare-disks".						
1.11.0 node_A_1	894.0GB	11	0	SSD	shared	testaggr
1.11.1 node_A_1	894.0GB	11	1	SSD	shared	testaggr
1.11.2 node_A_1	894.0GB	11	2	SSD	shared	testaggr

- If the original systems were configured with partitioned disks for ADP, enable it at the `LOADER` prompt for each replacement node:

```
setenv bootarg.mcc.adp_enabled true
```

- Set the following variables:

```
setenv bootarg.mcc.local_config_id <original-sys-id>
```

```
setenv bootarg.mcc.dr_partner <dr-partner-sys-id>
```



The `setenv bootarg.mcc.local_config_id` variable must be set to the sys-id of the **original** controller module, `node_B_1-old`.

- Set the variables on `node_B_1-new`.

The following example shows the commands for setting the values on `node_B_1-new`:

```
setenv bootarg.mcc.local_config_id 537403322
setenv bootarg.mcc.dr_partner 537403324
```

- b. Set the variables on node_B_2-new.

The following example shows the commands for setting the values on node_B_2-new:

```
setenv bootarg.mcc.local_config_id 537403321
setenv bootarg.mcc.dr_partner 537403323
```

6. If using encryption with external key manager, set the required bootargs:

```
setenv bootarg.kmip.init.ipaddr
setenv bootarg.kmip.kmip.init.netmask
setenv bootarg.kmip.kmip.init.gateway
setenv bootarg.kmip.kmip.init.interface
```

What's next?

[Reassign the root aggregate disks.](#)

Reassign the root aggregate disks to the new MetroCluster IP controller module

Reassign the root aggregate disks to the new controller module using the system IDs that you gathered earlier.

About this task

The old system IDs were identified in [Gather information before the upgrade](#).

You perform the steps in Maintenance mode.



Root aggregate disks are the only disks that must be reassigned during the controller upgrade process. Disk ownership of data aggregates is handled as part of the switchover/switchover operation.

Steps

1. Boot the system to Maintenance mode:

```
boot_ontap maint
```

2. Display the disks on node_B_1-new from the Maintenance mode prompt:

```
disk show -a
```



Before you proceed with disk reassignment, verify that the pool0 and pool1 disks that belong to the node's root aggregate are displayed in the `disk show` output. In the following example, the output lists the pool0 and pool1 disks owned by node_B_1-old.

The command output shows the system ID of the new controller module (1574774970). However, the old system ID (537403322) still owns the root aggregate disks. This example doesn't show drives owned by other nodes in the MetroCluster configuration.

```

*> disk show -a
Local System ID: 1574774970
DISK                OWNER                POOL  SERIAL NUMBER  HOME
DR HOME
-----
-----
prod3-rk18:9.126L44  node_B_1-old(537403322)  Pool1  PZHYN0MD
node_B_1-old(537403322)  node_B_1-old(537403322)
prod4-rk18:9.126L49  node_B_1-old(537403322)  Pool1  PPG3J5HA
node_B_1-old(537403322)  node_B_1-old(537403322)
prod4-rk18:8.126L21  node_B_1-old(537403322)  Pool1  PZHTDSZD
node_B_1-old(537403322)  node_B_1-old(537403322)
prod2-rk18:8.126L2   node_B_1-old(537403322)  Pool10  S0M1J2CF
node_B_1-old(537403322)  node_B_1-old(537403322)
prod2-rk18:8.126L3   node_B_1-old(537403322)  Pool10  S0M0CQM5
node_B_1-old(537403322)  node_B_1-old(537403322)
prod1-rk18:9.126L27  node_B_1-old(537403322)  Pool10  S0M1PSDW
node_B_1-old(537403322)  node_B_1-old(537403322)
.
.
.

```

3. Reassign the root aggregate disks on the drive shelves to the new controllers.

If you are using ADP...	Then use this command...
Yes	<code>disk reassign -s <old-sysid> -d <new-sysid> -r <dr-partner-sysid></code>
No	<code>disk reassign -s <old-sysid> -d <new-sysid></code>

4. Reassign the root aggregate disks on the drive shelves to the new controllers:

```
disk reassign -s <old-sysid> -d <new-sysid>
```

The following example shows reassignment of drives in a non-ADP configuration:

```
*> disk reassign -s 537403322 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 537403322.
Do you want to continue (y/n)? y
```

5. Verify that the disks of the root aggregate are correctly reassigned:

```
disk show
```

```
storage aggr status
```

```

*> disk show
Local System ID: 537097247

    DISK                                OWNER                                POOL  SERIAL NUMBER
HOME                                DR HOME
-----                                -
prod03-rk18:8.126L18 node_B_1-new(537097247) Pool1  PZHYN0MD
node_B_1-new(537097247) node_B_1-new(537097247)
prod04-rk18:9.126L49 node_B_1-new(537097247) Pool1  PPG3J5HA
node_B_1-new(537097247) node_B_1-new(537097247)
prod04-rk18:8.126L21 node_B_1-new(537097247) Pool1  PZHTDSZD
node_B_1-new(537097247) node_B_1-new(537097247)
prod02-rk18:8.126L2  node_B_1-new(537097247) Pool10 S0M1J2CF
node_B_1-new(537097247) node_B_1-new(537097247)
prod02-rk18:9.126L29 node_B_1-new(537097247) Pool10 S0M0CQM5
node_B_1-new(537097247) node_B_1-new(537097247)
prod01-rk18:8.126L1  node_B_1-new(537097247) Pool10 S0M1PSDW
node_B_1-new(537097247) node_B_1-new(537097247)
::>
::> aggr status
          Aggr              State              Status              Options
aggr0_node_B_1            online            raid_dp, aggr      root,
nosnap=on,
mirrored
mirror_resync_priority=high(fixed)
fast zeroed
64-bit

```

What's next?

[Boot the new controllers and restore LIF configuration.](#)

Boot the new MetroCluster IP controllers and restore LIF configuration

Boot the new controllers and verify that LIFs are hosted on appropriate nodes and ports.

Boot the new controllers

You must boot the new controllers, taking care to ensure that the bootarg variables are correct and, if needed, perform the encryption recovery steps.

Steps

1. Halt the new nodes:

```
halt
```

2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr <ip-address>
setenv bootarg.kmip.init.netmask <netmask>
setenv bootarg.kmip.init.gateway <gateway-address>
setenv bootarg.kmip.init.interface <interface-id>
```

3. Check if the partner-sysid is the current:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid <partner-sysID>
```

4. Display the ONTAP boot menu:

```
boot_ontap menu
```

5. If root encryption is used, select the boot menu option for your key management configuration.

If you are using...	Select this boot menu option...
Onboard key management	Option 10 Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.
External key management	Option 11 Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.

6. From the boot menu, select “(6) Update flash from backup config”.



Option 6 reboots the node twice before the process completes.

Respond with “y” to the system ID change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...
Rebooting to load the restored env file...
```

7. At the LOADER prompt, verify the bootarg values and update the values as needed.

Use the steps in [Set the MetroCluster IP bootarg variables](#).

8. Verify that the partner-sysid is the correct:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid <partner-sysID>
```

9. If root encryption is used, select the boot menu option again for your key management configuration.

If you are using...	Select this boot menu option...
Onboard key management	Option 10 Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.
External key management	Option "11" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.

Depending on the key manager setting, perform the recovery procedure by selecting option "10" or option "11", followed by option 6 at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option "1" (normal boot).

10. Wait for the replaced nodes to boot.

If either node is in takeover mode, perform a giveback using the `storage failover giveback` command.

11. If encryption is used, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> For more information, see Restore onboard key management encryption keys .
External key management	<pre>security key-manager external restore -vserver <SVM> -node <node> -key -server <host_name IP_address:port> -key-id key_id -key-tag key_tag <node_name></pre> For more information, see Restore external key management encryption keys .

12. Verify that all ports are in a broadcast domain:

a. View the broadcast domains:

```
network port broadcast-domain show
```

- b. If a new broadcast domain is created for the data ports on the newly upgraded controllers, delete the broadcast domain:



Only delete the new broadcast domain. Don't delete any of the broadcast domains that existed before starting the upgrade.

```
broadcast-domain delete -broadcast-domain <broadcast_domain_name>
```

- c. Add ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

- d. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might differ from the old node.

[Create a VLAN](#)

[Combine physical ports to create interface groups](#)

Verify and restore LIF configuration

Verify that LIFs are hosted on appropriate nodes and ports as mapped out at the beginning of the upgrade procedure.

About this task

- This task is performed on site_B.
- See the port mapping plan you created in [Map ports from the old nodes to the new nodes](#).



You must verify that the location of the data LIFs is correct on the new nodes before you perform a switchback. When you switchback the configuration, ONTAP attempts to resume traffic on the home port used by the LIFs. I/O failure can occur when the home port connection to the switch port and VLAN is incorrect.

Steps

1. Verify that LIFs are hosted on the appropriate node and ports before switchback.
 - a. Change to the advanced privilege level:

```
set -privilege advanced
```

- b. Display the LIFs, and confirm that each data LIF is using the correct home port:

```
network interface show
```

- c. Modify any LIFs that aren't using the correct home port:

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-port <port-id>
```

If the command returns an error, you can override the port configuration:

```
vserver config override -command "network interface modify -vserver <svm-name> -home-port <active_port_after_upgrade> -lif <lif_name> -home-node <new_node_name>"
```

When entering the network interface modify command within the `vserver config override` command, you cannot use the tab autocomplete feature. You can create the network interface modify using autocomplete and then enclose it in the `vserver config override` command.

- d. Confirm that all data LIFs are now on the correct home port:

```
network interface show
```

- e. Return to the admin privilege level:

```
set -privilege admin
```

2. Revert the interfaces to their home node:

```
network interface revert * -vserver <svm-name>
```

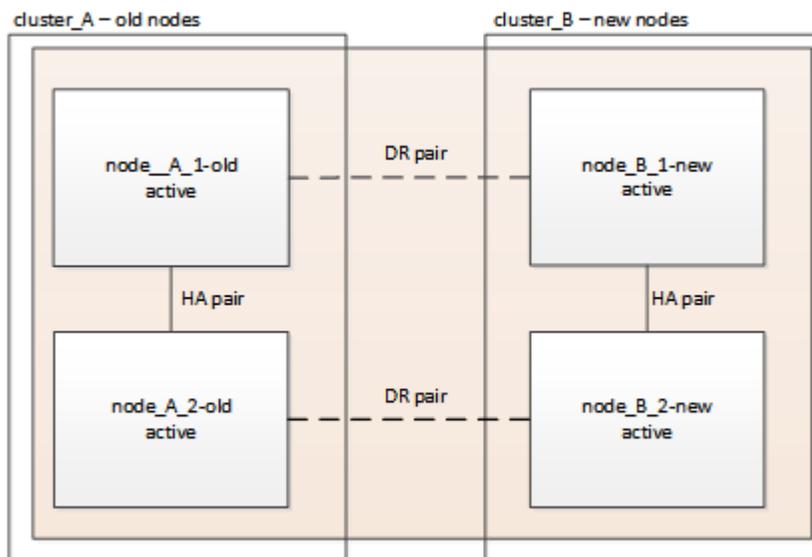
Perform this step on all SVMs as required.

What's next?

[Switchback the MetroCluster configuration.](#)

Switch back the MetroCluster IP configuration

Perform the switchback operation to return the MetroCluster configuration to normal operation. The nodes on site_A are still awaiting upgrade.



Steps

1. Issue the `metrocluster node show` command on site_B and check the output.
 - a. Verify that the new nodes are represented correctly.
 - b. Verify that the new nodes are in "Waiting for switchback state."

2. Perform the healing and switchback by running the required commands from any node in the active cluster (the cluster that is not undergoing upgrade).

a. Heal the data aggregates:

```
metrocluster heal aggregates
```

b. Heal the root aggregates:

```
metrocluster heal root
```

c. Switchback the cluster:

```
metrocluster switchback
```

3. Check the progress of the switchback operation:

```
metrocluster show
```

The switchback operation is still in progress when the output displays `waiting-for-switchback`:

```
cluster_B::> metrocluster show
Cluster                Entry Name              State
-----
Local: cluster_B      Configuration state    configured
                      Mode                    switchover
                      AUSO Failure Domain   -
Remote: cluster_A     Configuration state    configured
                      Mode                    waiting-for-switchback
                      AUSO Failure Domain   -
```

The switchback operation is complete when the output displays `normal`:

```
cluster_B::> metrocluster show
Cluster                Entry Name              State
-----
Local: cluster_B      Configuration state    configured
                      Mode                    normal
                      AUSO Failure Domain   -
Remote: cluster_A     Configuration state    configured
                      Mode                    normal
                      AUSO Failure Domain   -
```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command. This command is at the advanced privilege level.

What's next?

[Complete the upgrade.](#)

Complete the MetroCluster IP controller upgrade

After upgrading the controller modules, perform the necessary tasks to complete the controller upgrade.

Check the health of the MetroCluster configuration

After upgrading the controller modules you must verify the health of the MetroCluster configuration.

About this task

You can perform this task on any node in the MetroCluster configuration.

Steps

1. Verify the operation of the MetroCluster configuration:
 - a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Perform a MetroCluster check:

```
metrocluster check run
```

- c. Display the results of the MetroCluster check:

```
metrocluster check show
```

2. Verify the MetroCluster connectivity and status.

- a. Check the MetroCluster IP connections:

```
storage iscsi-initiator show
```

- b. Check that the nodes are operating:

```
metrocluster node show
```

- c. Check that the MetroCluster IP interfaces are up:

```
metrocluster configuration-settings interface show
```

- d. Check that local failover is enabled:

```
storage failover show
```

Upgrade the nodes on cluster_A

You must repeat the upgrade tasks to upgrade the nodes on cluster_A at site A.

Steps

1. Repeat the steps to upgrade the nodes on cluster_A, beginning with [Prepare for the upgrade.](#)

When you repeat the procedure, all example references to the clusters and nodes are reversed. For example, when the example is given to switchover from cluster_A, you will switchover from cluster_B.

Re-add the internal drives to the new controller

If you upgraded from a system that only has external drives to a system that has external and internal drives (disks and controllers in the same chassis), you can add or re-seat the disks that you removed or unseated from the internal slots of the new system. You can do this at any time after the upgrade is completed on both sites and the cluster is in a healthy state.

After you re-add or re-seat the drives, they can be used in ONTAP as required.



This task only applies to certain upgrade combinations. Refer to [remove internal drives from the chassis on the new controller](#) for more information.

Restore Tiebreaker or Mediator monitoring

After completing the upgrade of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

Steps

1. Restore monitoring if necessary, using the procedure for your configuration.

If you are using...	Use this procedure
Tiebreaker	Add MetroCluster configurations.
Mediator	Configure ONTAP Mediator from a MetroCluster IP configuration.
Third-party applications	Refer to the product documentation.

Send a custom AutoSupport message after maintenance

After completing the upgrade, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

Steps

1. To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.
 - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

Configure end-to-end encryption

If it is supported on your system, you can encrypt back-end traffic, such as NVlog and storage replication data, between the MetroCluster IP sites. Refer to [Configure end-to-end encryption](#) for more information.

Upgrade controllers from AFF A700/FAS9000 to AFF A900/FAS9500 in a MetroCluster IP configuration using switchover and switchback (ONTAP 9.10.1 or later)

You can use the MetroCluster switchover operation to provide nondisruptive service to clients while the controller modules on the partner cluster are upgraded. Other components (such as storage shelves or switches) cannot be upgraded as part of this procedure.

About this task

- To upgrade AFF A700 controller modules to AFF A900, the controllers must be running ONTAP 9.10.1 or later.
- To upgrade FAS9000 controller modules to FAS9500, the controllers must be running ONTAP 9.10.1P3 or later.
- All controllers in the configuration should be upgraded during the same maintenance period.

Operating the MetroCluster configuration with an AFF A700 and an AFF A900, or a FAS9000 and a FAS9500 controller is not supported outside of this maintenance activity.

- The IP switches must be running a supported firmware version.
- You will reuse the IP addresses, netmasks, and gateways of the original platforms on the new platforms.
- The following example names are used in this procedure, in both examples and graphics:
 - Site_A
 - Before upgrade:
 - node_A_1-A700
 - node_A_2-A700
 - After upgrade:
 - node_A_1-A900
 - node_A_2-A900
 - Site_B
 - Before upgrade:
 - node_B_1-A700
 - node_B_2-A700
 - After upgrade:
 - node_B_1-A900
 - node_B_2-A900

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.

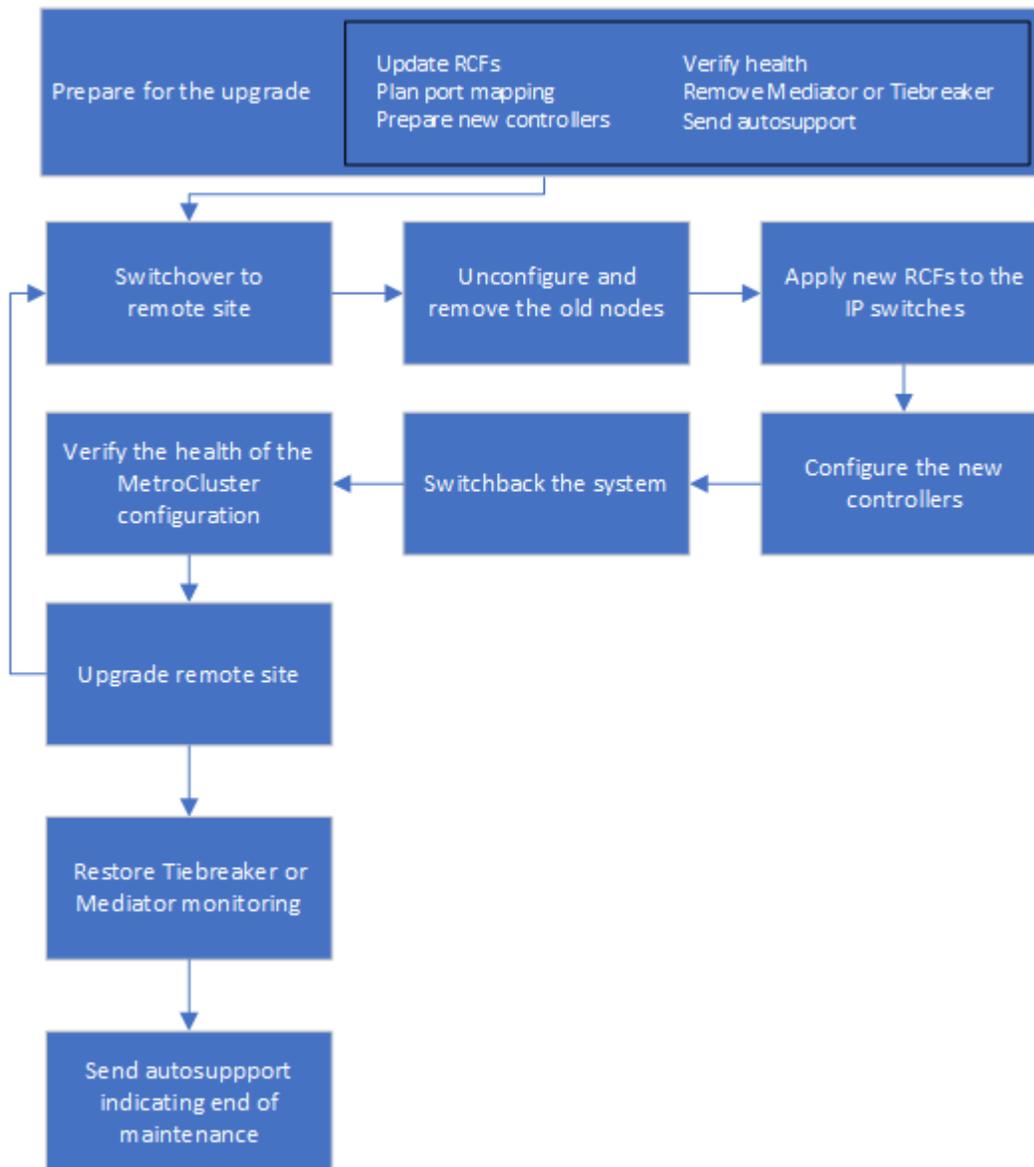
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Workflow for upgrading controllers in a MetroCluster IP configuration

You can use the workflow diagram to help you plan the upgrade tasks.



Prepare for the upgrade

Before making any changes to the existing MetroCluster configuration, you must check the health of the configuration, prepare the new platforms, and perform other miscellaneous tasks.

Clear slot 7 on the AFF A700 or FAS9000 controller

The MetroCluster configuration on an AFF A900 or FAS9500 uses one of each of the ports on the DR cards located in slots 5 and 7. Before starting the upgrade, if there are cards in slot 7 on the AFF A700 or FAS9000, you must move them to other slots for all the nodes of the cluster.

Update the MetroCluster switch RCF files before upgrading controllers

You must update the RCF files on MetroCluster switches when performing this upgrade. The following table provides the VLAN ranges supported for AFF A900/FAS9500 MetroCluster IP configurations.

Platform model	Supported VLAN IDs
<ul style="list-style-type: none">AFF A900 or FAS9500	<ul style="list-style-type: none">1020Any value in the range 101 to 4096 inclusive.

- If the switch is not configured with the minimum supported RCF file version, you must update the RCF file. For the correct RCF file version for your switch model, refer to the [RcfFileGenerator Tool](#). The following steps are for the RCF file application.

Steps

1. Prepare the IP switches for the application of the new RCF files.

Follow the steps in the section for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP switch to factory defaults](#)

2. Download and install the RCF files.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom RCF files](#)
- [Download and install the Cisco IP RCF files](#)
- [Download and install the NVIDIA IP RCF files](#)

Map ports from the old nodes to the new nodes

When upgrading from an AFF A700 to an AFF A900, or FAS9000 to FAS9500, you do not change the data network ports, FCP SAN adapter ports, and SAS and NVMe storage ports. Data LIFs stay where they are during and after the upgrade. Therefore, you are not required to map the network ports from the old nodes to the new nodes.

Verify MetroCluster health before site upgrade

You verify the health and connectivity of the MetroCluster configuration before performing the upgrade.



After you upgrade the controllers at the first site and before you upgrade the second, running `metrocluster check run` followed by `metrocluster check show` returns an error in the `config-replication` field. This error indicates an NVRAM size mismatch between the nodes at each site and it's the expected behavior when there are different platform models on both sites. You can ignore the error until the controller upgrade is completed for all nodes in the DR group.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

a. Check whether the nodes are multipathed:

```
node run -node node-name sysconfig -a
```

You should issue this command for each node in the MetroCluster configuration.

b. Verify that there are no broken disks in the configuration:

```
storage disk show -broken
```

You should issue this command on each node in the MetroCluster configuration.

c. Check for any health alerts:

```
system health alert show
```

You should issue this command on each cluster.

d. Verify the licenses on the clusters:

```
system license show
```

You should issue this command on each cluster.

e. Verify the devices connected to the nodes:

```
network device-discovery show
```

You should issue this command on each cluster.

f. Verify that the time zone and time is set correctly on both sites:

```
cluster date show
```

You should issue this command on each cluster. You can use the `cluster date` command to configure the time and time zone.

2. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

a. Confirm the MetroCluster configuration and that the operational mode is `normal`:

```
metrocluster show
```

b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

- c. Issue the following command:

```
metrocluster check run
```

- d. Display the results of the MetroCluster check:

```
metrocluster check show
```

3. Check the MetroCluster cabling with the Config Advisor tool.

- a. Download and run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- b. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Gather information before the upgrade

Before upgrading, you must gather information for each of the nodes, and, if necessary, adjust the network broadcast domains, remove any VLANs and interface groups, and gather encryption information.

Steps

1. Record the physical cabling for each node, labelling cables as needed to allow correct cabling of the new nodes.

2. Gather the output of the following commands for each node:

- `metrocluster interconnect show`
- `metrocluster configuration-settings connection show`
- `network interface show -role cluster,node-mgmt`
- `network port show -node node_name -type physical`
- `network port vlan show -node node-name`
- `network port ifgrp show -node node_name -instance`
- `network port broadcast-domain show`
- `network port reachability show -detail`
- `network ipspace show`
- `volume show`
- `storage aggregate show`
- `system node run -node node-name sysconfig -a`
- `vserver fcp initiator show`
- `storage disk show`
- `metrocluster configuration-settings interface show`

3. Gather the UUIDs for the site_B (the site whose platforms are currently being upgraded): `metrocluster node show -fields node-cluster-uuid, node-uuid`

These values must be configured accurately on the new site_B controller modules to ensure a successful upgrade. Copy the values to a file so that you can copy them into the proper commands later in the upgrade process.

The following example shows the command output with the UUIDs:

```
cluster_B::> metrocluster node show -fields node-cluster-uuid, node-uuid
(metrocluster node show)
dr-group-id cluster      node      node-uuid
node-cluster-uuid
-----
-----
1          cluster_A node_A_1-A700 f03cb63c-9a7e-11e7-b68b-00a098908039
ee7db9d5-9a82-11e7-b68b-00a098908039
1          cluster_A node_A_2-A700 aa9a7a7a-9a81-11e7-a4e9-00a098908c35
ee7db9d5-9a82-11e7-b68b-00a098908039
1          cluster_B node_B_1-A700 f37b240b-9ac1-11e7-9b42-00a098c9e55d
07958819-9ac6-11e7-9b42-00a098c9e55d
1          cluster_B node_B_2-A700 bf8e3f8f-9ac4-11e7-bd4e-00a098ca379f
07958819-9ac6-11e7-9b42-00a098c9e55d
4 entries were displayed.
cluster_B::~*
```

It is recommended that you record the UUIDs into a table similar to the following.

Cluster or node	UUID
cluster_B	07958819-9ac6-11e7-9b42-00a098c9e55d
node_B_1-A700	f37b240b-9ac1-11e7-9b42-00a098c9e55d
node_B_2-A700	bf8e3f8f-9ac4-11e7-bd4e-00a098ca379f
cluster_A	ee7db9d5-9a82-11e7-b68b-00a098908039
node_A_1-A700	f03cb63c-9a7e-11e7-b68b-00a098908039
node_A_2-A700	aa9a7a7a-9a81-11e7-a4e9-00a098908c35

4. If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

You should gather the output of the following commands:

- fcp adapter show -instance
- fcp interface show -instance
- iscsi interface show

° ucadmin show

5. If the root volume is encrypted, collect and save the passphrase used for key-manager:

```
security key-manager backup show
```

6. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.

For additional information, see [Backing up onboard key management information manually](#).

- a. If Onboard Key Manager is configured:

```
security key-manager onboard show-backup
```

You will need the passphrase later in the upgrade procedure.

- b. If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
security key-manager key query
```

7. Gather the system IDs of the existing nodes:

```
metrocluster node show -fields node-systemid,ha-partner-systemid,dr-partner-
systemid,dr-auxiliary-systemid
```

The following output shows the reassigned drives.

```
::> metrocluster node show -fields node-systemid,ha-partner-systemid,dr-
partner-systemid,dr-auxiliary-systemid

dr-group-id cluster      node      node-systemid ha-partner-systemid dr-
partner-systemid dr-auxiliary-systemid
-----
1          cluster_A node_A_1-A700  537403324      537403323
537403321          537403322
1          cluster_A node_A_2-A700  537403323      537403324
537403322          537403321
1          cluster_B node_B_1-A700  537403322      537403321
537403323          537403324
1          cluster_B node_B_2-A700  537403321      537403322
537403324          537403323
4 entries were displayed.
```

Remove Mediator or Tiebreaker monitoring

Before the upgrading the platforms, you must remove monitoring if the MetroCluster configuration is monitored with the Tiebreaker or Mediator utility.

Steps

1. Collect the output for the following command:

```
storage iscsi-initiator show
```

2. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

If you are using...	Use this procedure...
Tiebreaker	Removing MetroCluster Configurations in the <i>MetroCluster Tiebreaker Installation and Configuration content</i>
Mediator	Issue the following command from the ONTAP prompt: <pre>metrocluster configuration-settings mediator remove</pre>
Third-party applications	Refer to the product documentation.

Send a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. Log in to the cluster.
2. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

The `maintenance-window-in-hours` parameter specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

3. Repeat these steps on the partner site.

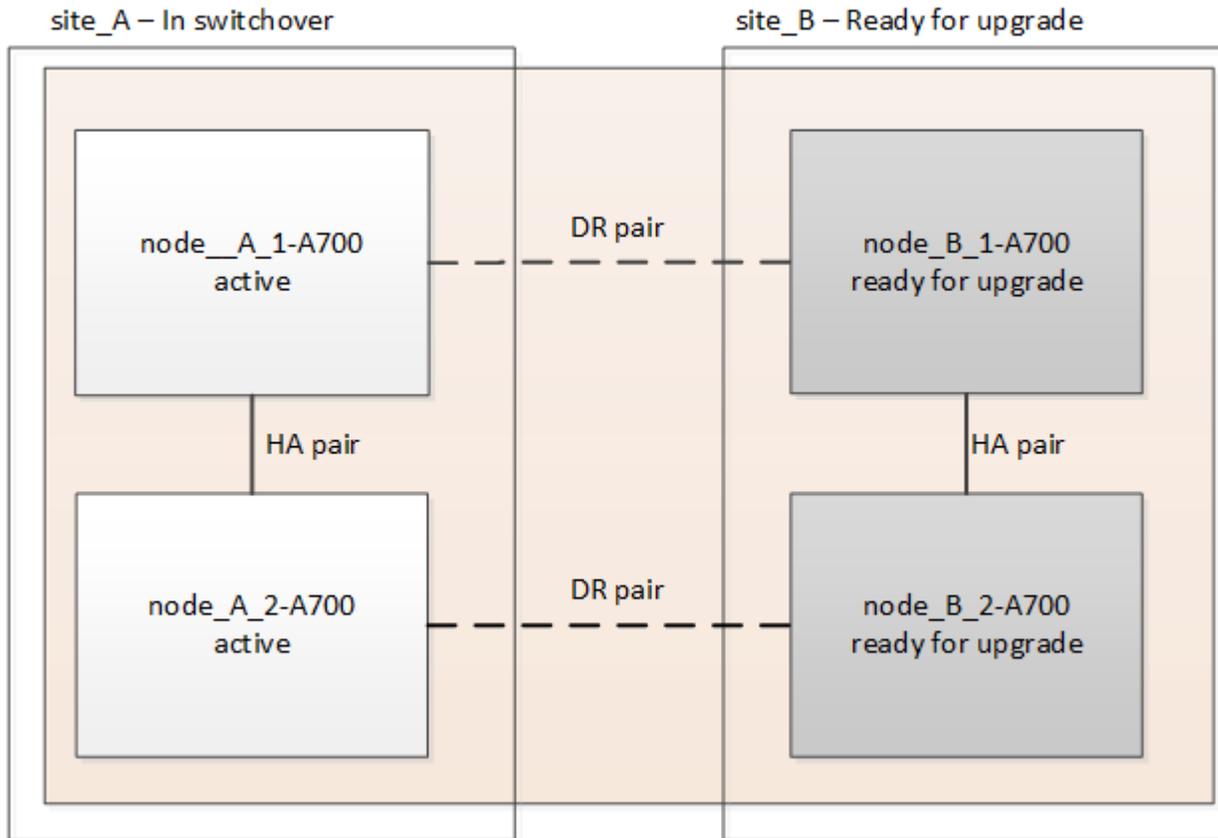
Switch over the MetroCluster configuration

You must switch over the configuration to `site_A` so that the platforms on `site_B` can be upgraded.

About this task

This task must be performed on `site_A`.

After completing this task, site_A is active and serving data for both sites. site_B is inactive, and ready to begin the upgrade process.



Steps

1. Switch over the MetroCluster configuration to site_A so that site_B's nodes can be upgraded:

a. Issue the following command on site_A:

```
metrocluster switchover -controller-replacement true
```

The operation can take several minutes to complete.

b. Monitor the switchover operation:

```
metrocluster operation show
```

c. After the operation is complete, confirm that the nodes are in switchover state:

```
metrocluster show
```

d. Check the status of the MetroCluster nodes:

```
metrocluster node show
```

Automatic healing of aggregates after negotiated switchover is disabled during controller upgrade. Nodes at site_B are halted and stopped at the `LOADER` prompt.

Remove AFF A700 or FAS9000 platform controller module and NVS

About this task

If you are not already grounded, properly ground yourself.

Steps

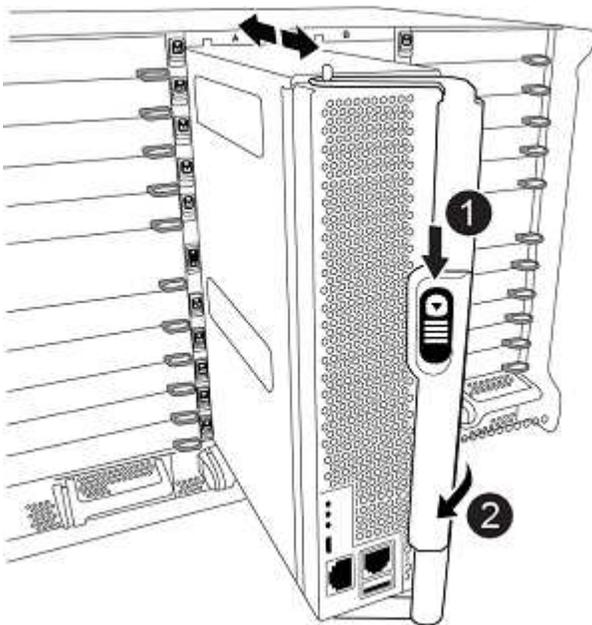
1. Gather the bootarg values from both nodes at site_B: `printenv`
2. Power off the chassis at site_B.

Remove the AFF A700 or FAS9000 controller module

Use the following procedure to remove the AFF A700 or FAS9000 controller module

Steps

1. Detach the console cable, if any, and the management cable from the controller module before removing the controller module.
2. Unlock and remove the controller module from the chassis.
 - a. Slide the orange button on the cam handle downward until it unlocks.



	Cam handle release button
	Cam handle

- b. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.
Make sure that you support the bottom of the controller module as you slide it out of the chassis.

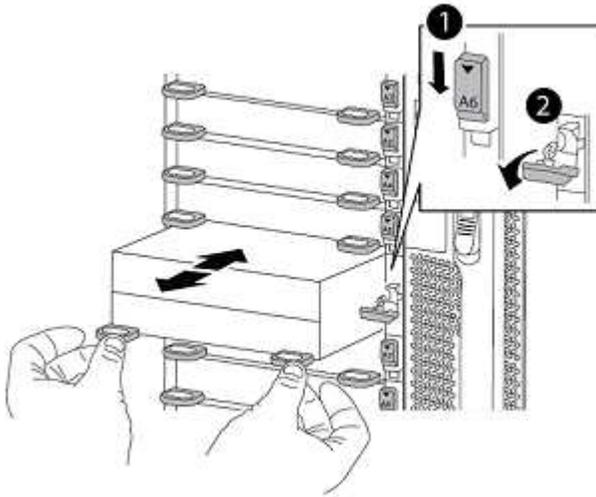
Remove the AFF A700 or FAS9000 NVS module

Use the following procedure to remove the AFF A700 or FAS9000 NVS module.

Note: The NVS module is in slot 6 and is double the height compared to other modules in the system.

Steps

1. Unlock and remove the NVS from slot 6.
 - a. Depress the lettered and numbered 'cam' button.
The cam button moves away from the chassis.
 - b. Rotate the cam latch down until it is in a horizontal position.
The NVS disengages from the chassis and moves a few inches.
 - c. Remove the NVS from the chassis by pulling on the pull tabs on the sides of the module face.



	Lettered and numbered I/O cam latch
	I/O latch completely unlocked

2. If you are using add-on modules used as coredump devices on the AFF A700 or FAS9000 NVS, do not transfer them to the AFF A900 or FAS9500 NVS.
Do not transfer any parts from the AFF A700 or FAS9000 controller module and NVS to the AFF A900 or FAS9500 module.

Install the AFF A900 or FAS9500 NVS and controller modules

You must install the AFF A900 or FAS9500 NVS and controller module that you received in the upgrade kit on both nodes at site_B. Do not move the coredump device from the AFF A700 or FAS9000 NVS module to the AFF A900 or FAS9500 NVS module.

About this task

If you are not already grounded, properly ground yourself.

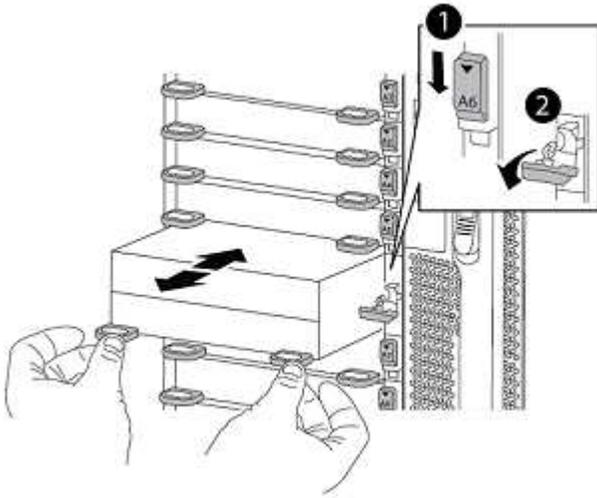
Install the AFF A900 or FAS9500 NVS

Use the following procedure to install the AFF A900 or FAS9500 NVS in slot 6 of both nodes at site_B.

Steps

1. Align the NVS with the edges of the chassis opening in slot 6.

2. Gently slide the NVS into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVS in place.



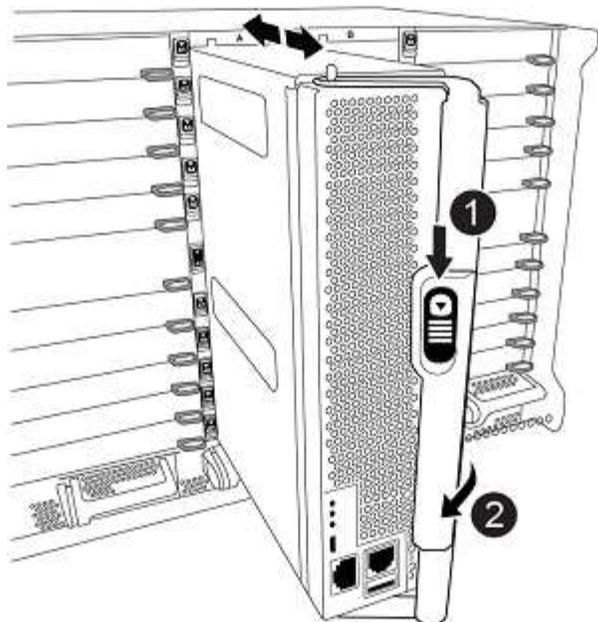
	Lettered and numbered I/O cam latch
	I/O latch completely unlocked

Install the AFF A900 or FAS9500 controller module.

Use the following procedure to install the AFF A900 or FAS9500 controller module.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Firmly push the controller module into the chassis until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated.
Attention: To avoid damaging the connectors, do not use excessive force when sliding the controller module into the chassis.
3. Cable the management and console ports to the controller module.



	Cam handle release button
	Cam handle

4. Install the second X91146A card in slot 7 of each node.
 - a. Move the e5b connection to e7b.
 - b. Move the e5a connection to e5b.



Slot 7 on all nodes of the cluster should be empty as mentioned in the [Map ports from the old nodes to the new nodes](#) section.

5. Power ON the chassis and connect to serial console.
6. After BIOS initialization, if the node starts autoboot, interrupt the AUTOBOOT by pressing Control-C.
7. After autoboot is interrupted, the nodes stop at the LOADER prompt. If you do not interrupt autoboot on time and node1 starts booting, wait for the prompt to press Ctrl-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt the autoboot during reboot.
8. At the LOADER prompt, set the default environment variables:
`set-defaults`
9. Save the default environment variables settings:
`saveenv`

Netboot nodes at site_B

After swapping the AFF A900 or FAS9500 controller module and NVS, you need to netboot the AFF A900 or FAS9500 nodes and install the same ONTAP version and patch level that is running on the cluster. The term netboot means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must add a copy of the ONTAP 9 boot image onto a web server that the system can access. It is not possible to check the version of ONTAP installed on the boot media of an AFF A900 or FAS9500 controller module unless it is installed in a chassis and powered ON. The ONTAP version on the AFF A900 or

FAS9500 boot media must be the same as the ONTAP version running on the AFF A700 or FAS9000 system that is being upgraded and both the primary and backup boot images should match. You can configure the images by performing a netboot followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the boot media.

Before you start

- Verify that you can access a HTTP server with the system.
- You need to download the necessary system files for your system and the correct version of ONTAP from the NetApp Support Site.

About this task

You must netboot the new controllers, if the version of ONTAP installed is not the same as the version installed on the original controllers. After you install each new controller, you boot the system from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots.

Steps

1. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `ontap-version_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available.
4. Your directory listing should contain `<ontap_version>_image.tgz`.
5. Configure the netboot connection by choosing one of the following actions.



You should use the management port and IP as the netboot connection. Do not use a data LIF IP or a data outage might occur while the upgrade is being performed.

If the Dynamic Host Configuration Protocol (DCHP) is...	Then...
Running	Configure the connection automatically by using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>

Not Running

Manually configure the connection by using the following command at the boot environment prompt:

```
ifconfig e0M -addr=<filer_addr>
-mask=<netmask> -gw=<gateway> -
dns=<dns_addr> domain=<dns_domain>
```

<filer_addr> is the IP address of the storage system. <netmask> is the network mask of the storage system.

<gateway> is the gateway for the storage system. <dns_addr> is the IP address of a name server on your network. This parameter is optional.

<dns_domain> is the Domain Name Service (DNS) domain name. This parameter is optional.

NOTE: Other parameters might be necessary for your interface. Enter `help ifconfig` at the firmware prompt for details.

6. Perform netboot on node_B_1:

```
netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel
```

The <path_to_the_web-accessible_directory> should lead to where you downloaded the <ontap_version>_image.tgz in [Step 2](#).



Do not interrupt the boot.

7. Wait for the node_B_1 now running on the AFF A900 or FAS9500 controller module to boot and display the boot menu options as shown below:

Please choose one of the following:

- (1) Normal Boot.
 - (2) Boot without /etc/rc.
 - (3) Change password.
 - (4) Clean configuration and initialize all disks.
 - (5) Maintenance mode boot.
 - (6) Update flash from backup config.
 - (7) Install new software first.
 - (8) Reboot node.
 - (9) Configure Advanced Drive Partitioning.
 - (10) Set Onboard Key Manager recovery secrets.
 - (11) Configure node for external key management.
- Selection (1-11)?

8. From the boot menu, select option (7) Install new software first.

This menu option downloads and installs the new ONTAP image to the boot device.

NOTE: Disregard the following message: This procedure is not supported for Non-Disruptive Upgrade on an HA pair. This note applies to nondisruptive ONTAP software

upgrades, and not controller upgrades.

Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the incorrect image might install. This issue applies to all ONTAP releases.

9. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

```
http://<web_server_ip/path_to_web-  
accessible_directory>/<ontap_version>\_image.tgz
```

10. Complete the following substeps to reboot the controller module:
 - a. Enter `n` to skip the backup recovery when you see the following prompt:

```
Do you want to restore the backup configuration now? {y|n} n
```

- b. Enter `y` to reboot when you see the following prompt:

```
The node must be rebooted to start using the newly installed  
software. Do you want to reboot now? {y|n} y
```

The controller module reboots but stops at the boot menu because the boot device was reformatted, and the configuration data needs to be restored.



You must reboot the node in order to use the newly installed software.

11. At the prompt, run the `wipeconfig` command to clear any previous configuration on the boot media:

- a. When you see the following message, answer `yes`:

```
This will delete critical system configuration, including cluster  
membership.
```

```
Warning: do not run this option on a HA node that has been taken over.  
Are you sure you want to continue?:
```

- b. The node reboots to finish the `wipeconfig` and then stops at the boot menu.

12. Select option 5 to go to maintenance mode from the boot menu. Answer `yes` to the prompts until the node stops at maintenance mode and the command prompt `*>`.

13. Repeat these steps to netboot node `_B_2`.

Restore the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site's usage.

Steps

1. In Maintenance mode configure the settings for any HBAs in the system:

- a. Check the current settings of the ports:

```
ucadmin show
```

b. Update the port settings as needed.

If you have this type of HBA and desired mode...	Use this command...
CNA FC	<code>ucadmin modify -m fc -t initiator adapter-name</code>
CNA Ethernet	<code>ucadmin modify -mode cna adapter-name</code>
FC target	<code>fcadmin config -t target adapter-name</code>
FC initiator	<code>fcadmin config -t initiator adapter-name</code>

2. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

3. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

```
boot_ontap maint
```

4. Verify the changes you made:

If you have this type of HBA...	Use this command...
CNA	<code>ucadmin show</code>
FC	<code>fcadmin show</code>

Set the HA state on the new controllers and chassis

You must verify the HA state of the controllers and chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be `mccip`.

2. If the displayed system state of the controller or chassis is not correct, set the HA state:

```
ha-config modify controller mccip
```

```
ha-config modify chassis mccip
```

3. Halt the node: `halt`

The node should stop at the `LOADER>` prompt.

4. On each node, check the system date, time, and time zone: `show date`

5. If necessary, set the date in UTC or GMT: `set date <mm/dd/yyyy>`

6. Check the time by using the following command at the boot environment prompt: `show time`

7. If necessary, set the time in UTC or GMT: `set time <hh:mm:ss>`

8. Save the settings: `saveenv`

9. Gather environment variables: `printenv`

Update the switch RCF files to accommodate the new platforms

You must update the switches to a configuration that supports the new platform models.

About this task

You perform this task at the site containing the controllers that are currently being upgraded. In the examples shown in this procedure we are upgrading `site_B` first.

The switches at `site_A` will be upgraded when the controllers on `site_A` are upgraded.

Steps

1. Prepare the IP switches for the application of the new RCFs.

Follow the steps in the section for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

2. Download and install the RCFs.

Follow the steps in the section for your switch vendor:

- [Download and install the Broadcom RCFs](#)
- [Download and install the Cisco IP RCFs](#)
- [Download and install the NVIDIA IP RCFs](#)

Configure the new controllers

New controllers should be ready and cabled at this point.

Set the MetroCluster IP bootarg variables

Certain MetroCluster IP bootarg values must be configured on the new controller modules. The values must match those configured on the old controller modules.

About this task

In this task, you will use the UUIDs and system IDs identified earlier in the upgrade procedure in [Gather](#)

information before the upgrade.

Steps

1. At the `LOADER>` prompt, set the following bootargs on the new nodes at site_B:

```
setenv bootarg.mcc.port_a_ip_config local-IP-address/local-IP-mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id
```

```
setenv bootarg.mcc.port_b_ip_config local-IP-address/local-IP-mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id
```

The following example sets the values for node_B_1-A900 using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12,120
setenv bootarg.mcc.port_b_ip_config
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12,130
```

The following example sets the values for node_B_2-A900 using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13,120
setenv bootarg.mcc.port_b_ip_config
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13,130
```

2. At the new nodes' `LOADER` prompt, set the UUIDs:

```
setenv bootarg.mgwd.partner_cluster_uuid partner-cluster-UUID
```

```
setenv bootarg.mgwd.cluster_uuid local-cluster-UUID
```

```
setenv bootarg.mcc.pri_partner_uuid DR-partner-node-UUID
```

```
setenv bootarg.mcc.aux_partner_uuid DR-aux-partner-node-UUID
```

```
setenv bootarg.mcc.iscsi.node_uuid local-node-UUID
```

- a. Set the UUIDs on node_B_1-A900.

The following example shows the commands for setting the UUIDs on node_B_1-A900:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.pri_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.aux_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f
setenv bootarg.mcc.iscsi.node_uuid f03cb63c-9a7e-11e7-b68b-
00a098908039
```

b. Set the UUIDs on node_B_2-A900:

The following example shows the commands for setting the UUIDs on node_B_2-A900:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.pri_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f
setenv bootarg.mcc.aux_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d
setenv bootarg.mcc.iscsi.node_uuid aa9a7a7a-9a81-11e7-a4e9-
00a098908c35
```

3. If the original systems were configured for ADP, at each of the replacement nodes' LOADER prompt, enable ADP:

```
setenv bootarg.mcc.adp_enabled true
```

4. Set the following variables:

```
setenv bootarg.mcc.local_config_id original-sys-id
```

```
setenv bootarg.mcc.dr_partner dr-partner-sys-id
```



The `setenv bootarg.mcc.local_config_id` variable must be set to the sys-id of the **original** controller module, node_B_1-A700.

a. Set the variables on node_B_1-A900.

The following example shows the commands for setting the values on node_B_1-A900:

```
setenv bootarg.mcc.local_config_id 537403322
setenv bootarg.mcc.dr_partner 537403324
```

- b. Set the variables on node_B_2-A900.

The following example shows the commands for setting the values on node_B_2-A900:

```
setenv bootarg.mcc.local_config_id 537403321
setenv bootarg.mcc.dr_partner 537403323
```

5. If using encryption with external key manager, set the required bootargs:

```
setenv bootarg.kmip.init.ipaddr
setenv bootarg.kmip.kmip.init.netmask
setenv bootarg.kmip.kmip.init.gateway
setenv bootarg.kmip.kmip.init.interface
```

Reassign root aggregate disks

Reassign the root aggregate disks to the new controller module, using the sysids gathered earlier.

About this task

These steps are performed in Maintenance mode.

Steps

1. Boot the system to Maintenance mode:

```
boot_ontap maint
```

2. Display the disks on node_B_1-A900 from the Maintenance mode prompt:

```
disk show -a
```

The command output shows the system ID of the new controller module (1574774970). However, the root aggregate disks are still owned by the old system ID (537403322). This example does not show drives owned by other nodes in the MetroCluster configuration.

```

*> disk show -a
Local System ID: 1574774970
DISK                OWNER                POOL  SERIAL NUMBER  HOME
DR HOME
-----
-----
prod3-rk18:9.126L44  node_B_1-A700(537403322)  Pool1  PZHYN0MD
node_B_1-A700(537403322)  node_B_1-A700(537403322)
prod4-rk18:9.126L49  node_B_1-A700(537403322)  Pool1  PPG3J5HA
node_B_1-A700(537403322)  node_B_1-700(537403322)
prod4-rk18:8.126L21  node_B_1-A700(537403322)  Pool1  PZHTDSZD
node_B_1-A700(537403322)  node_B_1-A700(537403322)
prod2-rk18:8.126L2   node_B_1-A700(537403322)  Pool10  SOM1J2CF
node_B_1-(537403322)  node_B_1-A700(537403322)
prod2-rk18:8.126L3   node_B_1-A700(537403322)  Pool10  SOM0CQM5
node_B_1-A700(537403322)  node_B_1-A700(537403322)
prod1-rk18:9.126L27  node_B_1-A700(537403322)  Pool10  SOM1PSDW
node_B_1-A700(537403322)  node_B_1-A700(537403322)
.
.
.

```

3. Reassign the root aggregate disks on the drive shelves to the new controllers.

If you are using ADP...	Then use this command...
Yes	<code>disk reassign -s <i>old-sysid</i> -d <i>new-sysid</i> -r <i>dr-partner-sysid</i></code>
No	<code>disk reassign -s <i>old-sysid</i> -d <i>new-sysid</i></code>

4. Reassign the root aggregate disks on the drive shelves to the new controllers:

```
disk reassign -s old-sysid -d new-sysid
```

The following example shows reassignment of drives in a non-ADP configuration:

```
*> disk reassign -s 537403322 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 537403322.
Do you want to continue (y/n)? y
```

5. Verify that the disks of the root aggregate are correctly reassigned old-remove:

```
disk show
```

```
storage aggr status
```

```

*> disk show
Local System ID: 537097247

    DISK                                OWNER                                POOL  SERIAL NUMBER
HOME                                DR HOME
-----                                -
-----                                -
prod03-rk18:8.126L18 node_B_1-A900 (537097247) Pool1  PZHYN0MD
node_B_1-A900 (537097247) node_B_1-A900 (537097247)
prod04-rk18:9.126L49 node_B_1-A900 (537097247) Pool1  PPG3J5HA
node_B_1-A900 (537097247) node_B_1-A900 (537097247)
prod04-rk18:8.126L21 node_B_1-A900 (537097247) Pool1  PZHTDSZD
node_B_1-A900 (537097247) node_B_1-A900 (537097247)
prod02-rk18:8.126L2  node_B_1-A900 (537097247) Pool10 SOM1J2CF
node_B_1-A900 (537097247) node_B_1-A900 (537097247)
prod02-rk18:9.126L29 node_B_1-A900 (537097247) Pool10 SOM0CQM5
node_B_1-A900 (537097247) node_B_1-A900 (537097247)
prod01-rk18:8.126L1  node_B_1-A900 (537097247) Pool10 SOM1PSDW
node_B_1-A900 (537097247) node_B_1-A900 (537097247)
::>
::> aggr status
          Aggr                State                Status                Options
aggr0_node_B_1                online                raid_dp, aggr                root,
nosnap=on,
                                mirrored
mirror_resync_priority=high(fixed)
                                fast zeroed
                                64-bit

```

Boot up the new controllers

You must boot the new controllers, taking care to ensure that the bootarg variables are correct and, if needed, perform the encryption recovery steps.

Steps

1. Halt the new nodes:

```
halt
```

2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address
```

```
setenv bootarg.kmip.init.netmask netmask
```

```
setenv bootarg.kmip.init.gateway gateway-address
```

```
setenv bootarg.kmip.init.interface interface-id
```

3. Check if the partner-sysid is the current:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid partner-sysID
```

4. Display the ONTAP boot menu:

```
boot_ontap menu
```

5. If root encryption is used, select the boot menu option for your key management configuration.

If you are using...	Select this boot menu option...
Onboard key management	Option 10 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration
External key management	Option 11 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration

6. From the boot menu, select (6) `Update flash from backup config`.



Option 6 will reboot the node twice before completing.

Respond `y` to the system id change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...  
  
Rebooting to load the restored env file...
```

7. Interrupt the AUTOBOOT to stop the controllers at LOADER.



On each node, check the bootargs set in [Setting the MetroCluster IP bootarg variables](#) and correct any incorrect values. Only move to the next step after you have checked the bootarg values.

8. Double-check that the partner-sysid is the correct:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid partner-sysID
```

9. If root encryption is used, select the boot menu option for your key management configuration.

If you are using...	Select this boot menu option...
Onboard key management	Option 10 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration
External key management	Option 11 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration

You need to perform the recovery procedure by selecting Option 10 or option 11 depending on the key manager setting and Option 6 at the boot menu prompt. To boot the nodes completely, you might need to perform the recovery procedure continued by option 1 (normal boot).

10. Wait for the new nodes, node_B_1-A900 and node_B_2-A900 to boot up.

If either node is in takeover mode, perform a giveback using the `storage failover giveback` command.

11. If encryption is used, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> <p>For more information, see Restoring onboard key management encryption keys.</p>
External key management	<pre>security key-manager external restore -vserver SVM -node node -key-server host_name IP_address:port -key-id key_id -key-tag key_tag node-name</pre> <p>For more information, see Restoring external key management encryption keys.</p>

12. Verify that all ports are in a broadcast domain:

a. View the broadcast domains:

```
network port broadcast-domain show
```

b. Add any ports to a broadcast domain as needed.

[Adding or removing ports from a broadcast domain](#)

c. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Creating a VLAN](#)

Verify and restore LIF configuration

Verify that LIFs are hosted on appropriate nodes and ports as mapped out at the beginning of the upgrade procedure.

About this task

- This task is performed on site_B.
- See the port mapping plan you created in [Map ports from the old nodes to the new nodes](#)

Steps

1. Verify that LIFs are hosted on the appropriate node and ports prior to switchback.

a. Change to the advanced privilege level:

```
set -privilege advanced
```

b. Override the port configuration to ensure proper LIF placement:

```
vserver config override -command "network interface modify -vserver  
vserver_name -home-port active_port_after_upgrade -lif lif_name -home-node  
new_node_name"
```

When entering the network interface modify command within the `vserver config override` command, you cannot use the tab autocomplete feature. You can create the network interface modify using autocomplete and then enclose it in the `vserver config override` command.

c. Return to the admin privilege level:

```
set -privilege admin
```

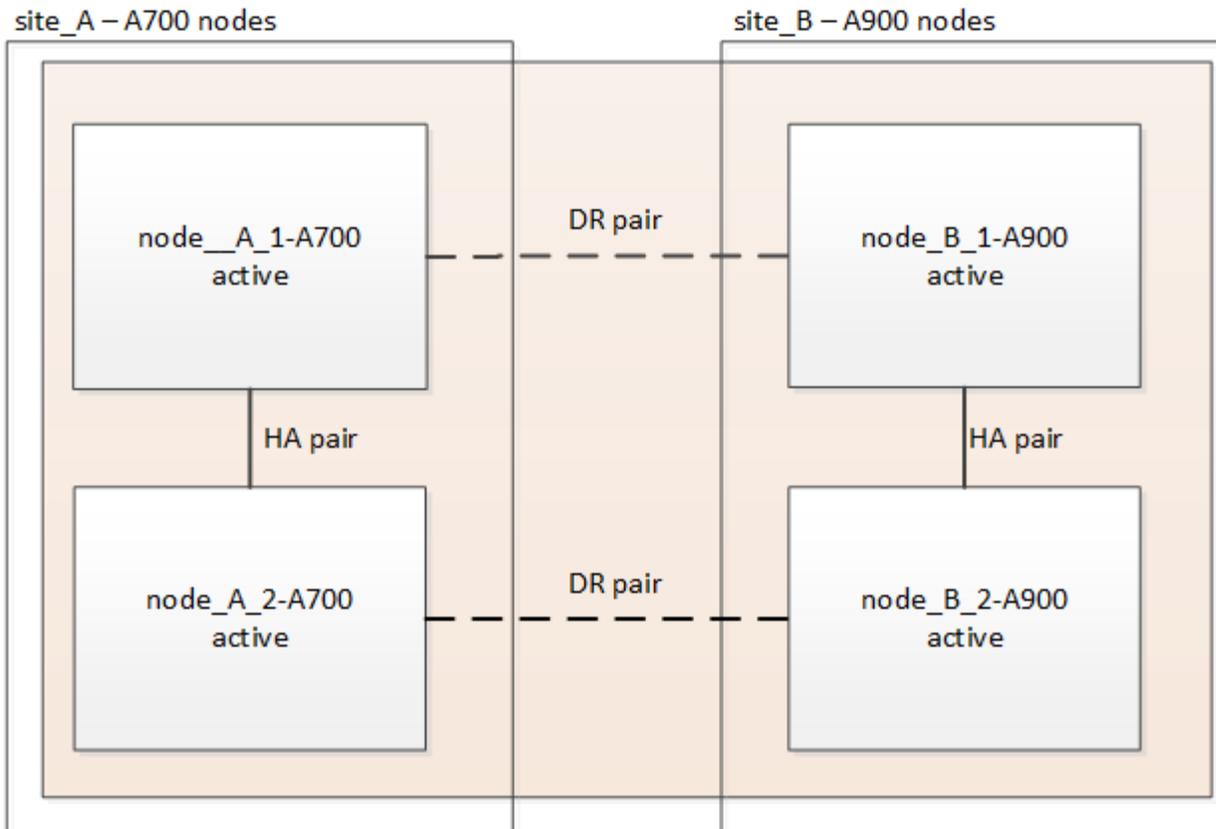
2. Revert the interfaces to their home node:

```
network interface revert * -vserver vserver-name
```

Perform this step on all SVMs as required.

Switch back the MetroCluster configuration

In this task, you will perform the switchback operation, and the MetroCluster configuration returns to normal operation. The nodes on site_A are still awaiting upgrade.



Steps

1. Issue the `metrocluster node show` command from `site_B` and check the output.
 - a. Verify that the new nodes are represented correctly.
 - b. Verify that the new nodes are in "Waiting for switchback state."
2. Perform the healing and switchback by running the required commands from any node in the active cluster (the cluster that is not undergoing upgrade).

- a. Heal the data aggregates:

```
metrocluster heal aggregates
```

- b. Heal the root aggregates:

```
metrocluster heal root
```

- c. Switchback the cluster:

```
metrocluster switchback
```

3. Check the progress of the switchback operation:

```
metrocluster show
```

The switchback operation is still in progress when the output displays `waiting-for-switchback`:

```

cluster_B::> metrocluster show
Cluster                Entry Name              State
-----
Local: cluster_B      Configuration state    configured
                      Mode                    switchover
                      AUSO Failure Domain -
Remote: cluster_A     Configuration state    configured
                      Mode                    waiting-for-switchback
                      AUSO Failure Domain -

```

The switchback operation is complete when the output displays normal:

```

cluster_B::> metrocluster show
Cluster                Entry Name              State
-----
Local: cluster_B      Configuration state    configured
                      Mode                    normal
                      AUSO Failure Domain -
Remote: cluster_A     Configuration state    configured
                      Mode                    normal
                      AUSO Failure Domain -

```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command. This command is at the advanced privilege level.

Check the health of the MetroCluster configuration

After upgrading the controller modules you must verify the health of the MetroCluster configuration.

About this task

This task can be performed on any node in the MetroCluster configuration.

Steps

1. Verify the operation of the MetroCluster configuration:
 - a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Perform a MetroCluster check:

```
metrocluster check run
```

- c. Display the results of the MetroCluster check:

```
metrocluster check show
```

2. Verify the MetroCluster connectivity and status.

a. Check the MetroCluster IP connections:

```
storage iscsi-initiator show
```

b. Check that the nodes are operating:

```
metrocluster node show
```

c. Check that the MetroCluster IP interfaces are up:

```
metrocluster configuration-settings interface show
```

d. Check that local failover is enabled:

```
storage failover show
```

Upgrade the nodes on site_A

You must repeat the upgrade tasks on site_A.

Steps

1. Repeat the steps to upgrade the nodes on site_A, beginning with [Prepare for the upgrade](#).

As you perform the tasks, all example references to the sites and nodes are reversed. For example, when the example is given to switchover from site_A, you will switchover from site_B.

Restore Tiebreaker or Mediator monitoring

After completing the upgrade of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

Steps

1. Restore monitoring if necessary, using the procedure for your configuration.

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations in the <i>MetroCluster Tiebreaker Installation and Configuration</i> section.
Mediator	Configure ONTAP Mediator from a MetroCluster IP configuration in the <i>MetroCluster IP Installation and Configuration</i> section.
Third-party applications	Refer to the product documentation.

Send a custom AutoSupport message after maintenance

After completing the upgrade, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

Steps

1. To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.
 - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

Upgrade controllers in a MetroCluster FC configuration using switchover and switchback

You can use the MetroCluster switchover operation to provide nondisruptive service to clients while the controller modules on the partner cluster are upgraded. Other components (such as storage shelves or switches) cannot be upgraded as part of this procedure.

Supported platform combinations

You can upgrade certain platforms using the switchover and switchback operation in a MetroCluster FC configuration.

For information on what platform upgrade combinations are supported review the MetroCluster FC upgrade table in [Choose a controller upgrade procedure](#).

Refer to [Choosing an upgrade or refresh method](#) for additional procedures.

About this task

- You can use this procedure only for controller upgrade.

Other components in the configuration, such as storage shelves or switches, cannot be upgraded at the same time.

- You can use this procedure with certain ONTAP versions:
 - Two-node configurations are supported in ONTAP 9.3 and later.
 - Four- and eight-node configurations are supported in ONTAP 9.8 and later.

Do not use this procedure on four- or eight-node configurations running ONTAP versions prior to 9.8.

- Your original and new platforms must be compatible and supported.

[NetApp Hardware Universe](#)



If the original or new platforms are FAS8020 or AFF8020 systems using ports 1c and 1d in FC-VI mode, see the Knowledge Base article [Upgrading controllers when FCVI connections on existing FAS8020 or AFF8020 nodes use ports 1c and 1d](#).

- The licenses at both sites must match. You can obtain new licenses from [NetApp Support](#).
- This procedure applies to controller modules in a MetroCluster FC configuration (a two-node stretch MetroCluster or a two, four-node, or eight-node fabric-attached MetroCluster configuration).
- All controllers in the same DR group should be upgraded during the same maintenance period.

Operating the MetroCluster configuration with different controller types in the same DR group is not supported outside of this maintenance activity. For eight-node MetroCluster configurations, the controllers within a DR Group must be the same, but both DR groups can use different controller types.

- Mapping of storage, FC and Ethernet connections between original nodes and new nodes in advance is recommended.
- If the new platform has fewer slots than the original system, or if it has fewer or different types of ports, you might need to add an adapter to the new system.

For more information, see the [NetApp Hardware Universe](#)

The following example names are used in this procedure:

- site_A
 - Before upgrade:
 - node_A_1-old
 - node_A_2-old
 - After upgrade:
 - node_A_1-new
 - node_A_2-new
- site_B
 - Before upgrade:
 - node_B_1-old
 - node_B_2-old
 - After upgrade:
 - node_B_1-new
 - node_B_2-new

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Prepare for the upgrade

Before making any changes to the existing MetroCluster configuration, you must check the health of the configuration, prepare the new platforms, and perform other miscellaneous tasks.

Verify the health of the MetroCluster configuration

You verify the health and connectivity of the MetroCluster configuration before performing the upgrade.



After you upgrade the controllers at the first site and before you upgrade the second, running `metrocluster check run` followed by `metrocluster check show` returns an error in the `config-replication` field. This error indicates an NVRAM size mismatch between the nodes at each site and it's the expected behavior when there are different platform models on both sites. You can ignore the error until the controller upgrade is completed for all nodes in the DR group.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

a. Check whether the nodes are multipathed:

```
node run -node node-name sysconfig -a
```

You should issue this command for each node in the MetroCluster configuration.

b. Verify that there are no broken disks in the configuration:

```
storage disk show -broken
```

You should issue this command on each node in the MetroCluster configuration.

c. Check for any health alerts:

```
system health alert show
```

You should issue this command on each cluster.

d. Verify the licenses on the clusters:

```
system license show
```

You should issue this command on each cluster.

e. Verify the devices connected to the nodes:

```
network device-discovery show
```

You should issue this command on each cluster.

- f. Verify that the time zone and time are set correctly on both sites:

```
cluster date show
```

You should issue this command on each cluster. You can use the `cluster date` commands to configure the time and time zone.

2. Check for any health alerts on the switches (if present):

```
storage switch show
```

You should issue this command on each cluster.

3. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

- a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

- c. Issue the following command:

```
metrocluster check run
```

- d. Display the results of the MetroCluster check:

```
metrocluster check show
```

4. Check the MetroCluster cabling with the Config Advisor tool.

- a. Download and run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- b. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Map ports from the old nodes to the new nodes

You must plan the mapping of the LIFs on physical ports on the old nodes to the physical ports on the new nodes.

About this task

When the new node is first booted during the upgrade process, it will replay the most recent configuration of the old node it is replacing. When you boot node_A_1-new, ONTAP attempts to host LIFs on the same ports that were used on node_A_1-old. Therefore, as part of the upgrade you must adjust the port and LIF configuration so it is compatible with that of the old node. During the upgrade procedure, you will perform steps on both the old and new nodes to ensure correct cluster, management, and data LIF configuration.

The following table shows examples of configuration changes related to the port requirements of the new nodes.

Cluster interconnect physical ports		
Old controller	New controller	Required action
e0a, e0b	e3a, e3b	No matching port. After the upgrade, recreate the cluster ports. Prepare cluster ports on an existing controller module
e0c, e0d	e0a,e0b,e0c,e0d	e0c and e0d are matching ports. You do not have to change the configuration, but after upgrade you can spread your cluster LIFs across the available cluster ports.

Steps

1. Determine what physical ports are available on the new controllers and what LIFs can be hosted on the ports.

The controller's port usage depends on the platform module and which switches you will use in the MetroCluster IP configuration. You can gather the port usage of the new platforms from the [NetApp Hardware Universe](#).

Also identify the FC-VI card slot usage.

2. Plan your port usage and, if desired, fill in the following tables for reference for each of the new nodes.

You will refer to the table as you carry out the upgrade procedure.

LIF	node_A_1-old			node_A_1-new		
	Ports	IPspaces	Broadcast domains	Ports	IPspaces	Broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						

Data 4						
SAN						
Intercluster port						

Gather information before the upgrade

Before upgrading, you must gather information for each of the old nodes, and, if necessary, adjust the network broadcast domains, remove any VLANs and interface groups, and gather encryption information.

About this task

This task is performed on the existing MetroCluster FC configuration.

Steps

1. Label the cables for the existing controllers, to allow easy identification of cables when setting up the new controllers.
2. Gather the system IDs of the nodes in the MetroCluster configuration:

```
metrocluster node show -fields node-systemid,dr-partner-systemid
```

During the upgrade procedure, you will replace these old system IDs with the system IDs of the new controller modules.

In this example for a four-node MetroCluster FC configuration, the following old system IDs are retrieved:

- node_A_1-old: 4068741258
- node_A_2-old: 4068741260
- node_B_1-old: 4068741254
- node_B_2-old: 4068741256

```

metrocluster-siteA::> metrocluster node show -fields node-
systemid,ha-partner-systemid,dr-partner-systemid,dr-auxiliary-
systemid
dr-group-id   cluster                               node
node-systemid   ha-partner-systemid   dr-partner-systemid
dr-auxiliary-systemid
-----
-----
-----
1               Cluster_A                               Node_A_1-old
4068741258      4068741260                               4068741256
4068741256
1               Cluster_A                               Node_A_2-old
4068741260      4068741258                               4068741254
4068741254
1               Cluster_B                               Node_B_1-old
4068741254      4068741256                               4068741258
4068741260
1               Cluster_B                               Node_B_2-old
4068741256      4068741254                               4068741260
4068741258
4 entries were displayed.

```

In this example for a two-node MetroCluster FC configuration, the following old system IDs are retrieved:

- node_A_1: 4068741258
- node_B_1: 4068741254

```

metrocluster node show -fields node-systemid,dr-partner-systemid
dr-group-id cluster   node           node-systemid dr-partner-systemid
-----
1           Cluster_A Node_A_1-old  4068741258    4068741254
1           Cluster_B node_B_1-old  -              -
2 entries were displayed.

```

3. Gather port and LIF information for each old node.

You should gather the output of the following commands for each node:

- network interface show -role cluster,node-mgmt
- network port show -node *node-name* -type physical
- network port vlan show -node *node-name*

- `network port ifgrp show -node node_name -instance`
- `network port broadcast-domain show`
- `network port reachability show -detail`
- `network ipspace show`
- `volume show`
- `storage aggregate show`
- `system node run -node node-name sysconfig -a`

4. If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

You should gather the output of the following commands:

- `fcg adapter show -instance`
- `fcg interface show -instance`
- `iscsi interface show`
- `ucadmin show`

5. If the root volume is encrypted, collect and save the passphrase used for key-manager:

```
security key-manager backup show
```

6. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.

For additional information, see [Backing up onboard key management information manually](#).

a. If Onboard Key Manager is configured:

```
security key-manager onboard show-backup
```

You will need the passphrase later in the upgrade procedure.

b. If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
```

```
security key-manager key query
```

Remove the existing configuration from the Tiebreaker or other monitoring software

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the Tiebreaker or other software prior to transition.

Steps

1. Remove the existing MetroCluster configuration from the Tiebreaker software.

[Remove MetroCluster configurations](#)

2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

Send a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway.
 - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

maintenance-window-in-hours specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

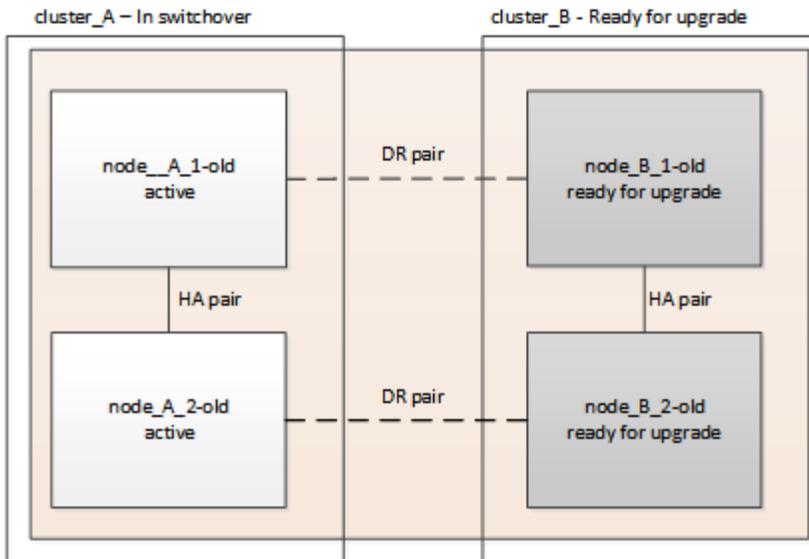
Switch over the MetroCluster configuration

You must switch over the configuration to site_A so that the platforms on site_B can be upgraded.

About this task

This task must be performed on site_A.

After completing this task, cluster_A is active and serving data for both sites. cluster_B is inactive, and ready to begin the upgrade process, as shown in the following illustration.



Steps

1. Switch over the MetroCluster configuration to site_A so that site_B's nodes can be upgraded:
 - a. Select the option that matches your configuration and issue the correct command on cluster_A:

Option 1: Four- or eight-node FC configuration running ONTAP 9.8 or later

Run the command: `metrocluster switchover -controller-replacement true`

Option 2: Two-node FC configuration running ONTAP 9.3 and later

Run the command: `metrocluster switchover`

The operation can take several minutes to complete.

- b. Monitor the switchover operation:

```
metrocluster operation show
```

- c. After the operation is complete, confirm that the nodes are in switchover state:

```
metrocluster show
```

- d. Check the status of the MetroCluster nodes:

```
metrocluster node show
```

2. Heal the data aggregates.

- a. Heal the data aggregates:

```
metrocluster heal data-aggregates
```

- b. Confirm the heal operation is complete by running the `metrocluster operation show` command on the healthy cluster:

```
cluster_A::> metrocluster operation show
  Operation: heal-aggregates
    State: successful
  Start Time: 7/29/2020 20:54:41
  End Time: 7/29/2020 20:54:42
  Errors: -
```

3. Heal the root aggregates.

a. Heal the data aggregates:

```
metrocluster heal root-aggregates
```

b. Confirm the heal operation is complete by running the `metrocluster operation show` command on the healthy cluster:

```
cluster_A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2020 20:58:41
  End Time: 7/29/2020 20:59:42
  Errors: -
```

Prepare the network configuration of the old controllers

To ensure that the networking resumes cleanly on the new controllers, you must move LIFs to a common port and then remove the networking configuration of the old controllers.

About this task

- This task must be performed on each of the old nodes.
- You will use the information gathered in [Mapping ports from the old nodes to the new nodes](#).

Steps

1. Boot the old nodes and then log in to the nodes:

```
boot_ontap
```

2. Assign the home port of all data LIFs on the old controller to a common port that is the same on both the old and new controller modules.

a. Display the LIFs:

```
network interface show
```

All data LIFS including SAN and NAS will be admin up and operationally down since those are up at switchover site (cluster_A).

b. Review the output to find a common physical network port that is the same on both the old and new

controllers that is not used as a cluster port.

For example, e0d is a physical port on old controllers and is also present on new controllers. e0d is not used as a cluster port or otherwise on the new controllers.

For port usage for platform models, see the [NetApp Hardware Universe](#)

- c. Modify all data LIFS to use the common port as the home port:

```
network interface modify -vserver svm-name -lif data-lif -home-port port-id
```

In the following example, this is "e0d".

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

3. Modify broadcast domains to remove vlan and physical ports that need to be deleted:

```
broadcast-domain remove-ports -broadcast-domain broadcast-domain-name -ports  
node-name:port-id
```

Repeat this step for all VLAN and physical ports.

4. Remove any VLAN ports using cluster ports as member ports and ifgrps using cluster ports as member ports.

- a. Delete VLAN ports:

```
network port vlan delete -node node-name -vlan-name portid-vlandid
```

For example:

```
network port vlan delete -node node1 -vlan-name e1c-80
```

- b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node node-name -ifgrp interface-group-name  
-port portid
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp ala -port e0d
```

- c. Remove VLAN and interface group ports from broadcast domain::

```
network port broadcast-domain remove-ports -ipSPACE ipSPACE -broadcast  
-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..
```

- d. Modify interface group ports to use other physical ports as member as needed.:

```
ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id
```

5. Halt the nodes:

```
halt -inhibit-takeover true -node node-name
```

This step must be performed on both nodes.

Remove the old platforms

The old controllers must be removed from the configuration.

About this task

This task is performed on site_B.

Steps

1. Connect to the serial console of the old controllers (node_B_1-old and node_B_2-old) at site_B and verify it is displaying the LOADER prompt.
2. Disconnect the storage and network connections on node_B_1-old and node_B_2-old and label the cables so they can be reconnected to the new nodes.
3. Disconnect the power cables from node_B_1-old and node_B_2-old.
4. Remove the node_B_1-old and node_B_2-old controllers from the rack.

Configure the new controllers

You must rack and install the controllers, perform required setup in Maintenance mode, and then boot the controllers, and verify the LIF configuration on the controllers.

Set up the new controllers

You must rack and cable the new controllers.

Steps

1. Plan out the positioning of the new controller modules and storage shelves as needed.

The rack space depends on the platform model of the controller modules, the switch types, and the number of storage shelves in your configuration.

2. Properly ground yourself.
3. Install the controller modules in the rack or cabinet.

[ONTAP Hardware Systems Documentation](#)

4. If the new controller modules did not come with FC-VI cards of their own and if FC-VI cards from old controllers are compatible on new controllers, swap FC-VI cards and install those in correct slots.

See the [NetApp Hardware Universe](#) for slot info for FC-VI cards.

5. Cable the controllers' power, serial console and management connections as described in the *MetroCluster Installation and Configuration Guides*.

Do not connect any other cables that were disconnected from old controllers at this time.

[ONTAP Hardware Systems Documentation](#)

6. Power up the new nodes and press Ctrl-C when prompted to display the LOADER prompt.

Netboot the new controllers

After you install the new nodes, you need to netboot to ensure the new nodes are running the same version of ONTAP as the original nodes. The term netboot means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

This task is performed on each of the new controller modules.

Steps

1. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `ontap-version_image.tgz` file on a web-accessible directory.
3. Go to the web-accessible directory and verify that the files you need are available.

If the platform model is...	Then...
FAS/AFF8000 series systems	<p>Extract the contents of the <code>ontap-version_image.tgz</code> file to the target directory: <code>tar -zxvf ontap-version_image.tgz</code></p> <p>NOTE: If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</p> <p>Your directory listing should contain a netboot folder with a kernel file: <code>netboot/kernel</code></p>
All other systems	<p>Your directory listing should contain a netboot folder with a kernel file: <code>ontap-version_image.tgz</code></p> <p>You do not need to extract the <code>ontap-version_image.tgz</code> file.</p>

4. At the LOADER prompt, configure the netboot connection for a management LIF:

- If IP addressing is DHCP, configure the automatic connection:

```
ifconfig e0M -auto
```

- If IP addressing is static, configure the manual connection:

```
ifconfig e0M -addr=ip_addr -mask=netmask -gw=gateway
```

5. Perform the netboot.

- If the platform is an 80xx series system, use this command:

```
netboot http://web_server_ip/path_to_web-accessible_directory/netboot/kernel
```

- If the platform is any other system, use the following command:

```
netboot http://web_server_ip/path_to_web-accessible_directory/ontap-  
version_image.tgz
```

6. From the boot menu, select option **(7) Install new software first** to download and install the new software image to the boot device.

```
Disregard the following message: "This procedure is not supported for  
Non-Disruptive Upgrade on an HA pair". It applies to nondisruptive  
upgrades of software, not to upgrades of controllers.
```

7. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL of the image file: `http://web_server_ip/path_to_web-accessible_directory/ontap-version_image.tgz`

```
Enter username/password if applicable, or press Enter to continue.
```

8. Enter `n` to skip the backup recovery when you see a prompt similar to the following:

```
Do you want to restore the backup configuration now? {y|n} n
```

9. Reboot by entering `y` when you see a prompt similar to the following:

```
The node must be rebooted to start using the newly installed software.  
Do you want to reboot now? {y|n} y
```



You must reboot the node in order to use the newly installed software.

Clear the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing configuration.

Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

Restore the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site's usage.

Steps

1. In Maintenance mode configure the settings for any HBAs in the system:

- a. Check the current settings of the ports: `ucadmin show`
- b. Update the port settings as needed.

If you have this type of HBA and desired mode...	Use this command...
CNA FC	<code>ucadmin modify -m fc -t initiator <i>adapter-name</i></code>
CNA Ethernet	<code>ucadmin modify -mode cna <i>adapter-name</i></code>
FC target	<code>fcadmin config -t target <i>adapter-name</i></code>
FC initiator	<code>fcadmin config -t initiator <i>adapter-name</i></code>

2. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the `LOADER` prompt.

3. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

```
boot_ontap maint
```

4. Verify the changes you made:

If you have this type of HBA...	Use this command...
CNA	<code>ucadmin show</code>
FC	<code>fcadmin show</code>

Set the HA state on the new controllers and chassis

You must verify the HA state of the controllers and chassis, and, if necessary, update the state to match your system configuration.

Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be `mcc`.

If the MetroCluster configuration has...	The HA state should be...
Two nodes	<code>mcc-2n</code>
Four or eight nodes	<code>mcc</code>

2. If the displayed system state of the controller is not correct, set the HA state for the controller module and chassis:

If the MetroCluster configuration has...	Issue these commands...
Two nodes	<code>ha-config modify controller mcc-2n</code> <code>ha-config modify chassis mcc-2n</code>
Four or eight nodes	<code>ha-config modify controller mcc</code> <code>ha-config modify chassis mcc</code>

Reassign root aggregate disks

Reassign the root aggregate disks to the new controller module, using the sysids gathered earlier

About this task

This task is performed in Maintenance mode.

The old system IDs were identified in [Gather information before the upgrade](#).

The examples in this procedure use controllers with the following system IDs:

Node	Old system ID	New system ID
node_B_1	4068741254	1574774970

Steps

1. Cable all other connections to the new controller modules (FC-VI, storage, cluster interconnect, etc.).
2. Halt the system and boot to Maintenance mode from the LOADER prompt:

```
boot_ontap maint
```

3. Display the disks owned by node_B_1-old:

```
disk show -a
```

The command output shows the system ID of the new controller module (1574774970). However, the root aggregate disks are still owned by the old system ID (4068741254). This example does not show drives owned by other nodes in the MetroCluster configuration.

```
*> disk show -a
Local System ID: 1574774970

   DISK          OWNER                                POOL  SERIAL NUMBER  HOME
DR HOME
-----
.....
...
rr18:9.126L44 node_B_1-old(4068741254)  Pool11 PZHYN0MD
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:9.126L49 node_B_1-old(4068741254)  Pool11 PPG3J5HA
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L21 node_B_1-old(4068741254)  Pool11 PZHTDSZD
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L2  node_B_1-old(4068741254)  Pool10 SOM1J2CF
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L3  node_B_1-old(4068741254)  Pool10 SOM0CQM5
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:9.126L27 node_B_1-old(4068741254)  Pool10 SOM1PSDW
node_B_1-old(4068741254) node_B_1-old(4068741254)
...

```

4. Reassign the root aggregate disks on the drive shelves to the new controller:

```
disk reassign -s old-sysid -d new-sysid
```

The following example shows reassignment of drives:

```

*> disk reassign -s 4068741254 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? Jul 14 19:23:49
[localhost:config.bridge.extra.port:error]: Both FC ports of FC-to-SAS
bridge rtp-fc02-41-rr18:9.126L0 S/N [FB7500N107692] are attached to this
controller.
y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 4068741254.
Do you want to continue (y/n)? y

```

5. Check that all disks are reassigned as expected:

```
disk show
```

```

*> disk show
Local System ID: 1574774970

  DISK          OWNER                                POOL   SERIAL NUMBER   HOME
DR HOME
-----
rr18:8.126L18  node_B_1-new(1574774970)   Pool1  PZHYN0MD
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L49  node_B_1-new(1574774970)   Pool1  PPG3J5HA
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L21  node_B_1-new(1574774970)   Pool1  PZHTDSZD
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L2   node_B_1-new(1574774970)   Pool0  SOM1J2CF
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L29  node_B_1-new(1574774970)   Pool0  SOM0CQM5
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L1   node_B_1-new(1574774970)   Pool0  SOM1PSDW
node_B_1-new(1574774970)  node_B_1-new(1574774970)
*>

```

6. Display the aggregate status:

```
aggr status
```

```
*> aggr status
      Aggr           State      Status      Options
aggr0_node_b_1-root  online    raid_dp, aggr  root, nosnap=on,
                    mirrored
mirror_resync_priority=high(fixed)
                    fast zeroed
                    64-bit
```

7. Repeat the above steps on the partner node (node_B_2-new).

Boot up the new controllers

You must reboot the controllers from the boot menu to update the controller flash image. Additional steps are required if encryption is configured.

About this task

This task must be performed on all the new controllers.

Steps

1. Halt the node:

```
halt
```

2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address
```

```
setenv bootarg.kmip.init.netmask netmask
```

```
setenv bootarg.kmip.init.gateway gateway-address
```

```
setenv bootarg.kmip.init.interface interface-id
```

3. Display the boot menu:

```
boot_ontap menu
```

4. If root encryption is used, depending on the ONTAP version you are using, select the boot menu option or issue the boot menu command for your key management configuration.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, select the boot menu option.

If you are using...	Select this boot menu option...
Onboard key management	Option "10" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.
External key management	Option "11" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, issue the boot menu command.

If you are using...	Issue this command at the boot menu prompt...
Onboard key management	<code>recover_onboard_keymanager</code>
External key management	<code>recover_external_keymanager</code>

5. If autoboot is enabled, interrupt autoboot by pressing CTRL-C.
6. From the boot menu, run option "6".



Option "6" will reboot the node twice before completing.

Respond "y" to the system id change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...  
  
Rebooting to load the restored env file...
```

7. Double-check that the partner-sysid is correct:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid partner-sysID
```

8. If root encryption is used, depending on the ONTAP version you are using, select the boot menu option or issue the boot menu command again for your key management configuration.

ONTAP 9.8 and later

Beginning with ONTAP 9.8, select the boot menu option.

If you are using...	Select this boot menu option...
Onboard key management	Option "10" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.
External key management	Option "11" Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.

Depending on the key manager setting, perform the recovery procedure by selecting option "10" or option "11", followed by option "6" at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option "1" (normal boot).

ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, issue the boot menu command.

If you are using...	Issue this command at the boot menu prompt...
Onboard key management	<code>recover_onboard_keymanager</code>
External key management	<code>recover_external_keymanager</code>

You might need to issue the `recover_XXXXXXX_keymanager` command at the boot menu prompt multiple times until the nodes completely boot.

9. Boot the nodes:

```
boot_ontap
```

10. Wait for the replaced nodes to boot up.

If either node is in takeover mode, perform a giveback:

```
storage failover giveback
```

11. Verify that all ports are in a broadcast domain:

- a. View the broadcast domains:

```
network port broadcast-domain show
```

- b. Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

- c. Add the physical port that will host the intercluster LIFs to the corresponding Broadcast domain.
- d. Modify intercluster LIFs to use the new physical port as home port.
- e. After the intercluster LIFs are up, check the cluster peer status and re-establish cluster peering as needed.

You might need to reconfigure cluster peering.

[Create a cluster peer relationship](#)

- f. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Create a VLAN](#)

[Combine physical ports to create interface groups](#)

12. If encryption is used, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> <p>For more information, see Restoring onboard key management encryption keys.</p>
External key management	<pre>security key-manager external restore -vserver SVM -node node -key-server host_name IP_address:port -key-id key_id -key-tag key_tag node-name</pre> <p>For more information, see Restoring external key management encryption keys.</p>

Verify LIF configuration

Verify that LIFs are hosted on appropriate node/ports prior to switchback. The following steps need to be performed

About this task

This task is performed on site_B, where the nodes have been booted up with root aggregates.

Steps

1. Verify that LIFs are hosted on the appropriate node and ports prior to switchback.
 - a. Change to the advanced privilege level:

```
set -privilege advanced
```

- b. Override the port configuration to ensure proper LIF placement:

```
vserver config override -command "network interface modify -vserver  
vserver_name -home-port active_port_after_upgrade -lif lif_name -home-node  
new_node_name"
```

When entering the `network interface modify` command within the `vserver config override` command, you cannot use the tab autocomplete feature. You can create the `network interface modify` using autocomplete and then enclose it in the `vserver config override` command.

- c. Return to the admin privilege level:

```
set -privilege admin
```

2. Revert the interfaces to their home node:

```
network interface revert * -vserver vserver-name
```

Perform this step on all SVMs as required.

Install the new licenses

Before the switchback operation, you must install licenses for the new controllers.

Steps

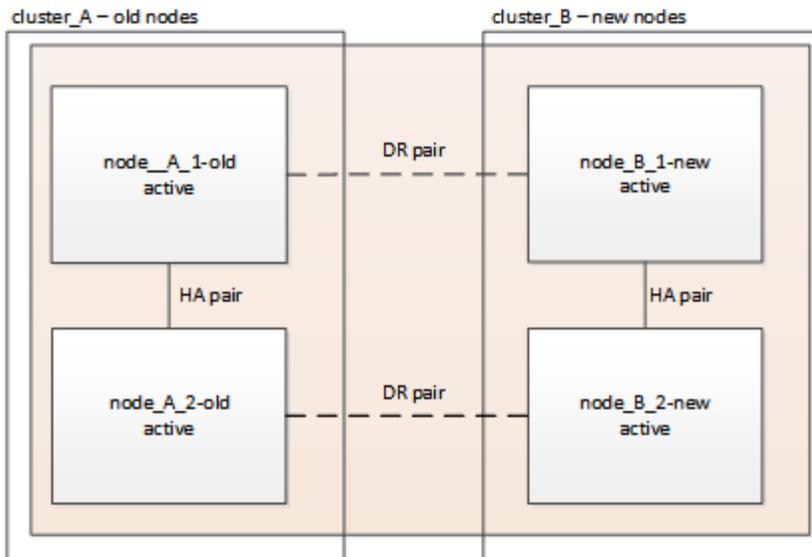
1. [Installing licenses for the new controller module](#)

Switch back the MetroCluster configuration

After the new controllers have been configured, you switch back the MetroCluster configuration to return the configuration to normal operation.

About this task

In this task, you will perform the switchback operation, returning the MetroCluster configuration to normal operation. The nodes on `site_A` are still awaiting upgrade.



Steps

1. Issue the `metrocluster node show` command on site_B and check the output.
 - a. Verify that the new nodes are represented correctly.
 - b. Verify that the new nodes are in "Waiting for switchback state."
2. Switchback the cluster:

```
metrocluster switchback
```

3. Check the progress of the switchback operation:

```
metrocluster show
```

The switchback operation is still in progress when the output displays `waiting-for-switchback`:

```
cluster_B::> metrocluster show
Cluster                Entry Name              State
-----
Local: cluster_B       Configuration state     configured
                       Mode                    switchover
                       AUSO Failure Domain    -
Remote: cluster_A     Configuration state     configured
                       Mode                    waiting-for-switchback
                       AUSO Failure Domain    -
```

The switchback operation is complete when the output displays `normal`:

```

cluster_B::> metrocluster show
Cluster                Entry Name                State
-----
Local: cluster_B      Configuration state      configured
                      Mode                      normal
                      AUSO Failure Domain     -
Remote: cluster_A     Configuration state      configured
                      Mode                      normal
                      AUSO Failure Domain     -

```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command. This command is at the advanced privilege level.

Check the health of the MetroCluster configuration

After upgrading the controller modules you must verify the health of the MetroCluster configuration.

About this task

This task can be performed on any node in the MetroCluster configuration.

Steps

1. Verify the operation of the MetroCluster configuration:
 - a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Perform a MetroCluster check:

```
metrocluster check run
```

- c. Display the results of the MetroCluster check:

```
metrocluster check show
```



After you run `metrocluster check run` and `metrocluster check show`, you see an error message similar to the following:

Example

```

Failed to validate the node and cluster components before the
switchover operation.
                Cluster_A:: node_A_1 (non-overridable veto): DR
partner NVLog mirroring is not online. Make sure that the links
between the two sites are healthy and properly configured.

```

This is expected behavior due to a controller mismatch during the upgrade process and the error

message can be safely ignored.

Upgrade the nodes on cluster_A

You must repeat the upgrade tasks on cluster_A.

Step

1. Repeat the steps to upgrade the nodes on cluster_A, beginning with [Prepare for the upgrade](#).

When you repeat the procedure, all example references to the clusters and nodes are reversed. For example, when the example is given to switchover from cluster_A, you will switchover from cluster_B.

Send a custom AutoSupport message after maintenance

After completing the upgrade, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

Step

1. To resume automatic support case generation, send an AutoSupport message to indicate that the maintenance is complete.
 - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

Restore Tiebreaker monitoring

If the MetroCluster configuration was previously configured for monitoring by the Tiebreaker software, you can restore the Tiebreaker connection.

1. Use the steps in [Add MetroCluster configurations](#) in *MetroCluster Tiebreaker Installation and Configuration*.

Upgrade controllers from AFF A700/FAS9000 to AFF A900/FAS9500 in a MetroCluster FC configuration using switchover and switchback (ONTAP 9.10.1 or later)

You can use the MetroCluster switchover operation to provide nondisruptive service to clients while the controller modules on the partner cluster are upgraded. You cannot upgrade other components (such as storage shelves or switches) as part of this procedure.

About this task

- You can use this procedure only for controller upgrade.

You cannot upgrade other components in the configuration, such as storage shelves or switches, at the same time.

- You can use this procedure to upgrade an AFF A700 to AFF A900 with ONTAP 9.10.1 and later.

- You can use this procedure to upgrade a FAS9000 to FAS9500 with ONTAP 9.10.1P3 and later.
 - Four and eight-node configurations are supported in ONTAP 9.10.1 and later.



The AFF A900 system is only supported in ONTAP 9.10.1 or later.

[NetApp Hardware Universe](#)

- All controllers in the configuration should be upgraded during the same maintenance period.

The following table shows the supported model matrix for the controller upgrade.

Old platform model	New platform model
<ul style="list-style-type: none"> • AFF A700 	<ul style="list-style-type: none"> • AFF A900
<ul style="list-style-type: none"> • FAS9000 	<ul style="list-style-type: none"> • FAS9500

- During the upgrade procedure, you are required to change the MetroCluster fabric, including the RCF and physical changes of cabling. You can perform the RCF and cabling changes before performing the controller upgrade.
- This upgrade procedure does not require you do not change the storage, FC, and Ethernet connections between the original nodes and the new nodes.
- During the upgrade procedure, you should not add or remove other cards from the AFF A700 or FAS9000 system. For more information, see the [NetApp Hardware Universe](#)

The following example names are used in examples and graphics in this procedure:

- site_A
 - Before upgrade:
 - node_A_1-A700
 - node_A_2-A700
 - After upgrade:
 - node_A_1-A900
 - node_A_2-A900
- site_B
 - Before upgrade:
 - node_B_1-A700
 - node_B_2-A700
 - After upgrade:
 - node_B_1-A900
 - node_B_2-A900

Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the

following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

Prepare for the upgrade

Before making any changes to the existing MetroCluster configuration, you must check the health of the configuration, change the RCF files and cabling to match to new port connectivity topology required for the AFF A900 or FAS9000 fabric MetroCluster configuration, and perform other miscellaneous tasks.

Clear slot 7 on the AFF A700 controller

The MetroCluster configuration on an AFF A900 or FAS9500 requires 8 FC-VI ports across FC-VI cards in slots 5 and 7. Before starting the upgrade, if there are cards in slot 7 on the AFF A700 or FAS9000, you must move them to other slots for all the nodes of the cluster.

Verify the health of the MetroCluster configuration

Before you update the RCF files and cabling for the AFF A900 or FAS9500 fabric MetroCluster configuration, you must verify the health and connectivity of the configuration.



After you upgrade the controllers at the first site and before you upgrade the second, running `metrocluster check run` followed by `metrocluster check show` returns an error in the `config-replication` field. This error indicates an NVRAM size mismatch between the nodes at each site and it's the expected behavior when there are different platform models on both sites. You can ignore the error until the controller upgrade is completed for all nodes in the DR group.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the nodes are multipathed:

```
node run -node node-name sysconfig -a
```

You should issue this command for each node in the MetroCluster configuration.

- b. Verify that there are no broken disks in the configuration:

```
storage disk show -broken
```

You should issue this command on each node in the MetroCluster configuration.

- c. Check for any health alerts:

```
system health alert show
```

You should issue this command on each cluster.

- d. Verify the licenses on the clusters:

```
system license show
```

You should issue this command on each cluster.

- e. Verify the devices connected to the nodes:

```
network device-discovery show
```

You should issue this command on each cluster.

- f. Verify that the time zone and time are set correctly on both sites:

```
cluster date show
```

You should issue this command on each cluster. You can use the `cluster date` commands to configure the time and time zone.

2. Check for any health alerts on the switches (if present):

```
storage switch show
```

You should issue this command on each cluster.

3. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

- a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

- c. Issue the following command:

```
metrocluster check run
```

- d. Display the results of the MetroCluster check:

```
metrocluster check show
```

4. Check the MetroCluster cabling with the Config Advisor tool.

- a. Download and run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- b. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Update the fabric switch RCF files

The AFF A900 or FAS9500 fabric MetroCluster requires two four-port FC-VI adapters per node compared to a single four-port FC-VI adapter required by an AFF A700. Before you start the controller upgrade to the AFF A900 or FAS9500 controller, you must modify the fabric switch RCF files to support the AFF A900 or FAS9500 connection topology.

1. From the [MetroCluster RCF file download page](#), download the correct RCF file for an AFF A900 or FAS9500 fabric MetroCluster and the switch model that is in use on the AFF A700 or FAS9000 configuration.
2. Update the RCF file on the fabric A switches, switch A1, and switch B1 by following the steps in [Configuring the FC switches](#).



The RCF file update to support the AFF A900 or FAS9500 fabric MetroCluster configuration does not affect the port and connections used for the AFF A700 or FAS9000 fabric MetroCluster configuration.

3. After updating the RCF files on the fabric A switches, all storage and FC-VI connections should come online. Check the FC-VI connections:

```
metrocluster interconnect mirror show
```

- a. Verify that the local and remote site disks are listed in the `sysconfig` output.
4. You must verify that MetroCluster is in a healthy state after the RCF file update for fabric A switches.
 - a. Check metro cluster connections:

```
metrocluster interconnect mirror show
```
 - b. Run metrocluster check:

```
metrocluster check run
```
 - c. See the MetroCluster run results when the run completes:

```
metrocluster check show
```
 5. Update the fabric B switches (switches 2 and 4) by repeating [Step 2](#) to [Step 5](#).

Verify the health of the MetroCluster configuration after the RCF file update

You must verify the health and connectivity of the MetroCluster configuration before performing the upgrade.

Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:
 - a. Check whether the nodes are multipathed:

```
node run -node node-name sysconfig -a
```

You should issue this command for each node in the MetroCluster configuration.

- b. Verify that there are no broken disks in the configuration:

```
storage disk show -broken
```

You should issue this command on each node in the MetroCluster configuration.

c. Check for any health alerts:

```
system health alert show
```

You should issue this command on each cluster.

d. Verify the licenses on the clusters:

```
system license show
```

You should issue this command on each cluster.

e. Verify the devices connected to the nodes:

```
network device-discovery show
```

You should issue this command on each cluster.

f. Verify that the time zone and time are set correctly on both sites:

```
cluster date show
```

You should issue this command on each cluster. You can use the `cluster date` commands to configure the time and time zone.

2. Check for any health alerts on the switches (if present):

```
storage switch show
```

You should issue this command on each cluster.

3. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

c. Issue the following command:

```
metrocluster check run
```

d. Display the results of the MetroCluster check:

```
metrocluster check show
```

4. Check the MetroCluster cabling with the Config Advisor tool.

a. Download and run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- b. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

Map ports from the AFF A700 or FAS9000 nodes to the AFF A900 or FAS9500 nodes

During the controller upgrade process, you must only change the connections that are mentioned in this procedure.

If the AFF A700 or FAS9000 controllers have a card in slot 7, you should move it to another slot before starting the controller upgrade procedure. You must have slot 7 available for the addition of the second FC-VI adapter that is required for the functioning of fabric MetroCluster on the AFF A900 or FAS9500 controllers.

Gather information before the upgrade

Before upgrading, you must gather information for each of the old nodes, and, if necessary, adjust the network broadcast domains, remove any VLANs and interface groups, and gather encryption information.

About this task

This task is performed on the existing MetroCluster FC configuration.

Steps

1. Gather the MetroCluster configuration node system IDs:

```
metrocluster node show -fields node-systemid,dr-partner-systemid
```

During the upgrade procedure, you will replace these old system IDs with the system IDs of the controller modules.

In this example for a four-node MetroCluster FC configuration, the following old system IDs are retrieved:

- node_A_1-A700: 537037649
- node_A_2-A700: 537407030
- node_B_1-A700: 0537407114
- node_B_2-A700: 537035354

```

Cluster_A::*> metrocluster node show -fields node-systemid,ha-partner-
systemid,dr-partner-systemid,dr-auxiliary-systemid
dr-group-id cluster      node          node-systemid ha-partner-systemid
dr-partner-systemid dr-auxiliary-systemid
-----
-----
1          Cluster_A  nodeA_1-A700  537407114      537035354
537411005          537410611
1          Cluster_A  nodeA_2-A700  537035354      537407114
537410611          537411005
1          Cluster_B  nodeB_1-A700  537410611      537411005
537035354          537407114
1          Cluster_B  nodeB_2-A700  537411005

4 entries were displayed.

```

2. Gather port and LIF information for each old node.

You should gather the output of the following commands for each node:

- network interface show -role cluster,node-mgmt
- network port show -node *node-name* -type physical
- network port vlan show -node *node-name*
- network port ifgrp show -node *node_name* -instance
- network port broadcast-domain show
- network port reachability show -detail
- network ipspace show
- volume show
- storage aggregate show
- system node run -node *node-name* sysconfig -a

3. If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

You should gather the output of the following commands:

- fcp adapter show -instance
- fcp interface show -instance
- iscsi interface show
- ucadmin show

4. If the root volume is encrypted, collect and save the passphrase used for key-manager:

```
security key-manager backup show
```

5. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.

For additional information, see [Backing up onboard key management information manually](#).

- a. If Onboard Key Manager is configured:

```
security key-manager onboard show-backup
```

You will need the passphrase later in the upgrade procedure.

- b. If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
```

```
security key-manager key query
```

Remove the existing configuration from the Tiebreaker or other monitoring software

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the Tiebreaker or other software prior to transition.

Steps

1. Remove the existing MetroCluster configuration from the Tiebreaker software.

[Removing MetroCluster Configurations](#)

2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

Send a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

1. To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway.

- a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-  
window-in-hours
```

`maintenance-window-in-hours` specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

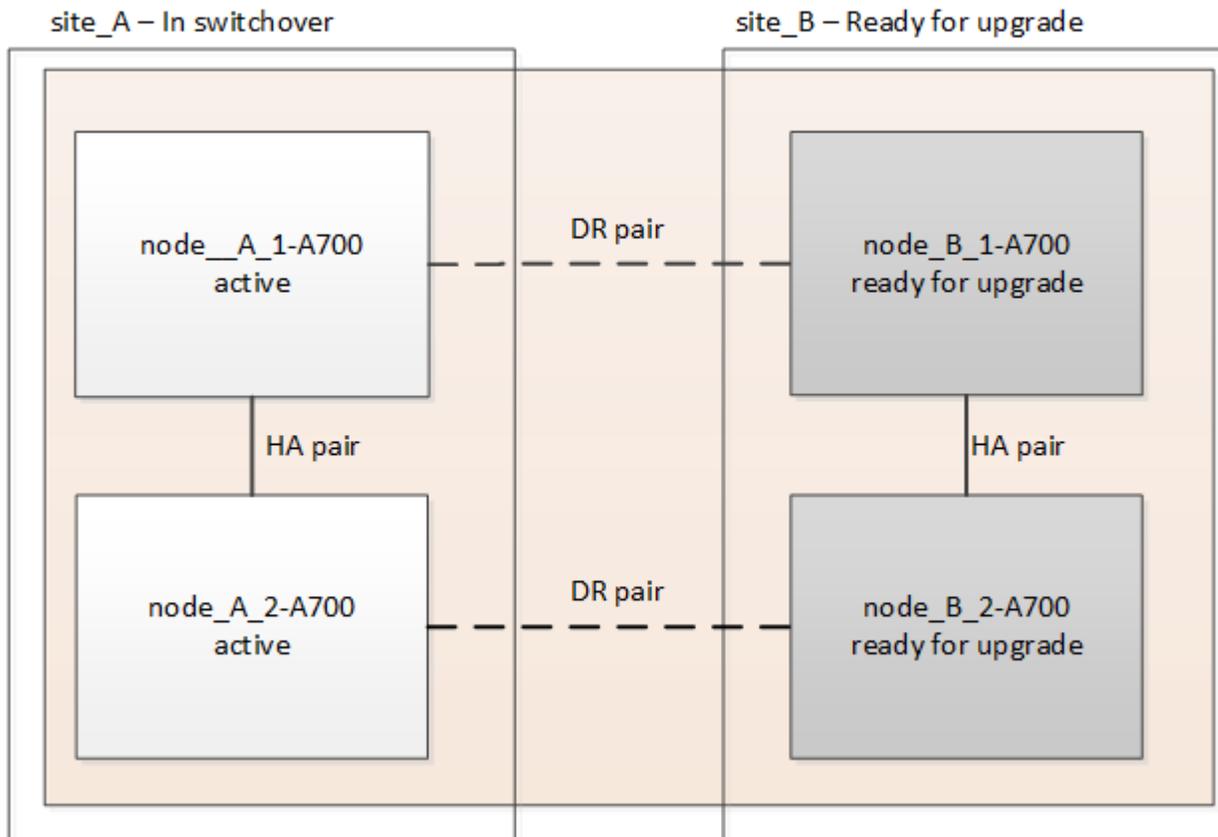
Switch over the MetroCluster configuration

You must switch over the configuration to site_A so that the platforms on site_B can be upgraded.

About this task

This task must be performed on site_A.

After completing this task, site_A is active and serving data for both sites. Site_B is inactive, and ready to begin the upgrade process, as shown in the following illustration. (This illustration also applies to upgrading a FAS9000 to a FAS9500 controller.)



Steps

1. Switch over the MetroCluster configuration to site_A so that site_B's nodes can be upgraded:

- a. Issue the following command on site_A:

```
metrocluster switchover -controller-replacement true
```

The operation can take several minutes to complete.

- b. Monitor the switchover operation:

```
metrocluster operation show
```

- c. After the operation is complete, confirm that the nodes are in switchover state:

```
metrocluster show
```

- d. Check the status of the MetroCluster nodes:

```
metrocluster node show
```

2. Heal the data aggregates.

- a. Heal the data aggregates:

```
metrocluster heal data-aggregates
```

- b. Confirm the heal operation is complete by running the `metrocluster operation show` command on the healthy cluster:

```
cluster_A::> metrocluster operation show
Operation: heal-aggregates
State: successful
Start Time: 7/29/2020 20:54:41
End Time: 7/29/2020 20:54:42
Errors: -
```

3. Heal the root aggregates.

- a. Heal the data aggregates:

```
metrocluster heal root-aggregates
```

- b. Confirm the heal operation is complete by running the `metrocluster operation show` command on the healthy cluster:

```
cluster_A::> metrocluster operation show
Operation: heal-root-aggregates
State: successful
Start Time: 7/29/2020 20:58:41
End Time: 7/29/2020 20:59:42
Errors: -
```

Remove the AFF A700 or FAS9000 controller module and NVS at site_B

You must remove the old controllers from the configuration.

You perform this task on site_B.

Before you begin

If you are not already grounded, properly ground yourself.

Steps

1. Connect to the serial console of the old controllers (node_B_1-700 and node_B_2-700) at site_B and verify it is displaying the `LOADER` prompt.
2. Gather the bootarg values from both nodes at site_B: `printenv`
3. Power off the chassis at site_B.

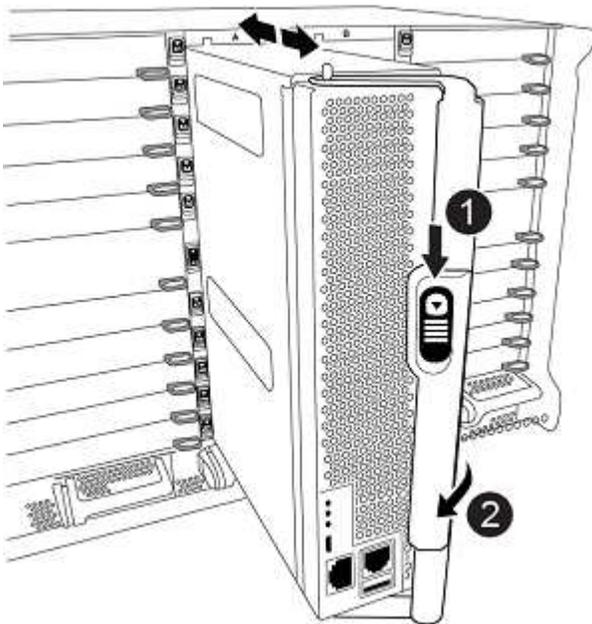
Remove the controller module and NVS from both nodes at site_B

Remove the AFF A700 or FAS9000 controller module

Use the following procedure to remove the AFF A700 or FAS9000 controller module.

Steps

1. Detach the console cable, if any, and the management cable from the controller module before removing the controller module.
2. Unlock and remove the controller module from the chassis.
 - a. Slide the orange button on the cam handle downward until it unlocks.



	Cam handle release button
	Cam handle

- b. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.
Make sure that you support the bottom of the controller module as you slide it out of the chassis.

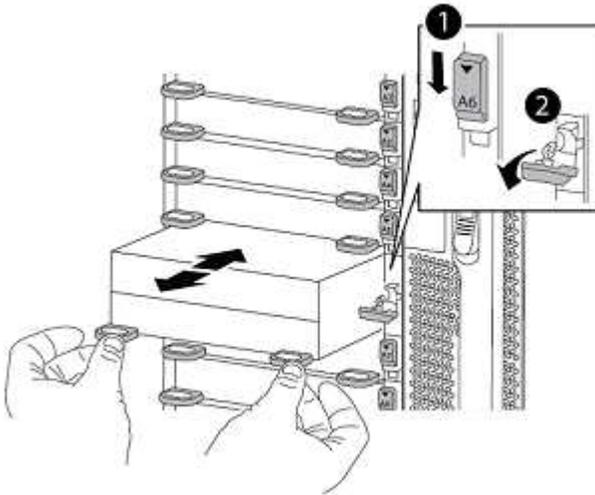
Remove the AFF A700 or FAS9000 NVS module

Use the following procedure to remove the AFF A700 or FAS9000 NVS module.



The AFF A700 or FAS9000 NVS module is in slot 6 and is double the height compared to the other modules in the system.

1. Unlock and remove the NVS from slot 6.
 - a. Depress the lettered and numbered cam button.
The cam button moves away from the chassis.
 - b. Rotate the cam latch down until it is in a horizontal position.
The NVS disengages from the chassis and moves a few inches.
 - c. Remove the NVS from the chassis by pulling on the pull tabs on the sides of the module face.



	Lettered and numbered I/O cam latch
	I/O latch completely unlocked



- Do not transfer any add-on modules used as coredump devices on the AFF A700 non-volatile storage module in slot 6 to the AFF A900 NVS module. Do not transfer any parts from the AFF A700 controller and NVS modules to the AFF A900 controller module.
- For FAS9000 to FAS9500 upgrades, you should only transfer Flash Cache modules on the FAS9000 NVS module to the FAS9500 NVS module. Do not transfer any other parts from the FAS9000 controller and NVS modules to the FAS9500 controller module.

Install the AFF A900 or FAS9500 NVS and controller module

You must install the AFF A900 or FAS9500 NVS and controller module from the upgrade kit on both nodes at Site_B. Do not move the coredump device from the AFF A700 or FAS9000 NVS module to the AFF A900 or FAS9500 NVS module.

Before you start

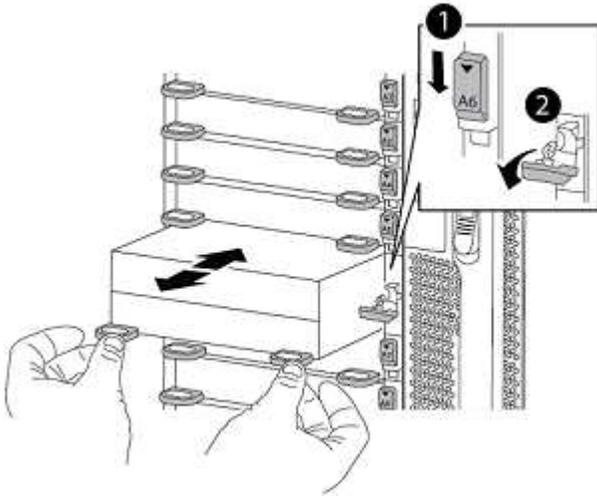
If you are not already grounded, properly ground yourself.

Install the AFF A900 or FAS9500 NVS

Use the following procedure to install the AFF A900 or FAS9500 NVS in slot 6 of both nodes at site_B

Steps

1. Align the NVS with the edges of the chassis opening in slot 6.
2. Gently slide the NVS into the slot until the lettered and numbered I/O cam latch begins to engage with the I/O cam pin, and then push the I/O cam latch all the way up to lock the NVS in place.



	Lettered and numbered I/O cam latch
	I/O latch completely unlocked

Install the AFF A900 or FAS9500 controller module

Use the following procedure to install the AFF A900 or FAS9500 controller module.

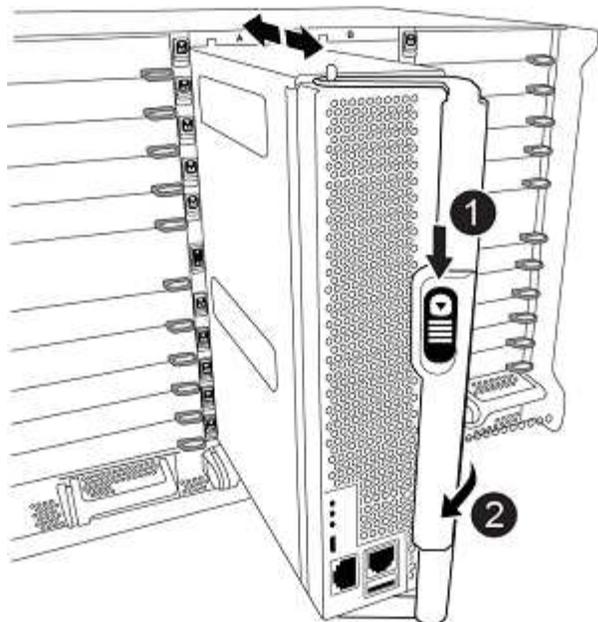
Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.
2. Firmly push the controller module into the chassis until it meets the midplane and is fully seated. The locking latch rises when the controller module is fully seated.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

3. Cable the management and console ports to the controller module.



	Cam handle release button
	Cam handle

4. Install the second X91129A card in slot 7 of each node.
 - a. Connect FC-VI ports from slot 7 to the switches. Refer to the [Fabric-attached installation and configuration](#) documentation and go to the AFF A900 or FAS9500 fabric MetroCluster connection requirements for the type of switch in your environment.
5. Power ON the chassis and connect to the serial console.
6. After BIOS initialization, if the node starts to autoboot, interrupt the AUTOBOOT by pressing Control-C.
7. After you interrupt the autoboot, the nodes stop at the LOADER prompt. If you do not interrupt autoboot on time and node1 starts booting, wait for the prompt to press Control-C to go into the boot menu. After the node stops at the boot menu, use option 8 to reboot the node and interrupt the autoboot during the reboot.
8. At the LOADER prompt, set the default environment variables: `set-defaults`
9. Save the default environment variables settings: `saveenv`

Netboot the nodes at site_B

After swapping the AFF A900 or FAS9500 controller module and NVS, you need to netboot the AFF A900 or FAS9500 nodes and install the same ONTAP version and patch level that is running on the cluster. The term `netboot` means you are booting from an ONTAP image stored on a remote server. When preparing for `netboot`, you must add a copy of the ONTAP 9 boot image onto a web server that the system can access.

It is not possible to check the ONTAP version installed on the boot media of an AFF A900 or FAS9500 controller module unless it is installed in a chassis and powered ON. The ONTAP version on the AFF A900 or FAS9500 boot media must be same as the ONTAP version running on the AFF A700 or FAS9000 system that is being upgraded and both the primary and backup boot images should match. You can configure the images by performing a `netboot` followed by the `wipeconfig` command from the boot menu. If the controller module was previously used in another cluster, the `wipeconfig` command clears any residual configuration on the

boot media.

Before you start

- Verify that you can access a HTTP server with the system.
- You need to download the necessary system files for your system and the correct version of ONTAP from the [NetApp Support](#) site.

About this task

You must `netboot` the new controllers if the version of ONTAP installed is not the same as the version installed on the original controllers. After you install each new controller, you boot the system from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots.

Steps

1. Access [NetApp Support](#) to download the files required to perform a system netboot used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
3. Change to the web-accessible directory and verify that the files you need are available. Your directory listing should contain `<ontap_version>_image.tgz`.
4. Configure the `netboot` connection by choosing one of the following actions.
Note: You should use the management port and IP as the `netboot` connection. Do not use a data LIF IP or a data outage might occur while the upgrade is being performed.

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Running	Configure the connection automatically by using the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>
Not running	Manually configure the connection by using the following command at the boot environment prompt: <code>ifconfig e0M -addr=<filer_addr> -mask=<netmask> -gw=<gateway> -dns=<dns_addr> domain=<dns_domain></code> <filer_addr> is the IP address of the storage system. <netmask> is the network mask of the storage system. <gateway> is the gateway for the storage system. <dns_addr> is the IP address of a name server on your network. This parameter is optional. <dns_domain> is the Domain Name Service (DNS) domain name. This parameter is optional. NOTE: Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.

5. Perform `netboot` on node 1:
`netboot http://<web_server_ip/path_to_web_accessible_directory>/netboot/kernel`

The `<path_to_the_web-accessible_directory>` should lead to where you downloaded the `<ontap_version>_image.tgz` in [Step 2](#).



Do not interrupt the boot.

6. Wait for node 1 that is running on the AFF A900 or FAS9500 controller module to boot and display the boot menu options as shown below:

```
Please choose one of the following:
```

- ```
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)?
```

7. From the boot menu, select option (7) Install new software first. This menu option downloads and installs the new ONTAP image to the boot device.



Disregard the following message: This procedure is not supported for Non-Disruptive Upgrade on an HA pair. This note applies to nondisruptive ONTAP software upgrades, and not controller upgrades.

Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the wrong incorrect image might install. This issue applies to all ONTAP releases.

8. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL:

```
http://<web_server_ip/path_to_web-
accessible_directory>/<ontap_version>_image.tgz
```

9. Complete the following substeps to reboot the controller module:

- a. Enter `n` to skip the backup recovery when you see the following prompt:

```
Do you want to restore the backup configuration now? {y|n} n
```

- b. Enter `y` to reboot when you see the following prompt:

The node must be rebooted to start using the newly installed software. Do you want to reboot now? {y|n} y

The controller module reboots but stops at the boot menu because the boot device was reformatted, and the configuration data needs to be restored.



You must reboot the node in order to use the newly installed software.

10. At the prompt, run the `wipeconfig` command to clear any previous configuration on the boot media:
  - a. When you see the message below, answer `yes`:

```
This will delete critical system configuration, including cluster membership.
Warning: do not run this option on a HA node that has been taken over.
Are you sure you want to continue?:
```
  - b. The node reboots to finish the `wipeconfig` and then stops at the boot menu.
11. Select option 5 to go to maintenance mode from the boot menu. Answer `yes` to the prompts until the node stops at maintenance mode and the command prompt `*>`.

## Restore the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site's usage.

### Steps

1. In Maintenance mode configure the settings for any HBAs in the system:
  - a. Check the current settings of the ports: `ucadmin show`
  - b. Update the port settings as needed.

| If you have this type of HBA and desired mode... | Use this command...                                         |
|--------------------------------------------------|-------------------------------------------------------------|
| CNA FC                                           | <code>ucadmin modify -m fc -t initiator adapter-name</code> |
| CNA Ethernet                                     | <code>ucadmin modify -mode cna adapter-name</code>          |
| FC target                                        | <code>fcadmin config -t target adapter-name</code>          |
| FC initiator                                     | <code>fcadmin config -t initiator adapter-name</code>       |

## Set the HA state on the new controllers and chassis

You must verify the HA state of the controllers and chassis, and, if necessary, update the state to match your system configuration.

### Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be `mcc`.

2. If the displayed system state of the controller or chassis is not correct, set the HA state:

```
ha-config modify controller mcc
```

```
ha-config modify chassis mcc
```

3. Halt the node: `halt`  
The node should stop at the `LOADER>` prompt.
4. On each node, check the system date, time, and time zone: `show date`
5. If necessary, set the date in UTC or Greenwich Mean Time (GMT): `set date <mm/dd/yyyy>`
6. Check the time by using the following command at the boot environment prompt: `show time`
7. If necessary, set the time in UTC or GMT: `set time <hh:mm:ss>`
8. Save the settings: `saveenv`
9. Gather environment variables: `printenv`
10. Boot the node back into Maintenance mode to enable the configuration changes to take effect:  
`boot_ontap maint`
11. Verify the changes you made are effective and `ucadmin` shows FC initiator ports online.

| If you have this type of HBA... | Use this command...       |
|---------------------------------|---------------------------|
| CNA                             | <code>ucadmin show</code> |
| FC                              | <code>fcadmin show</code> |

12. Verify the `ha-config` mode: `ha-config show`
  - a. Verify that you have the following output:

```
*> ha-config show
Chassis HA configuration: mcc
Controller HA configuration: mcc
```

### Set the HA state on the new controllers and chassis

You must verify the HA state of the controllers and chassis, and, if necessary, update the state to match your system configuration.

#### Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be mcc.

| If the MetroCluster configuration has... | The HA state should be... |
|------------------------------------------|---------------------------|
| Two nodes                                | mcc-2n                    |
| Four or eight nodes                      | mcc                       |

2. If the displayed system state of the controller is not correct, set the HA state for the controller module and chassis:

| If the MetroCluster configuration has... | Issue these commands...                                                           |
|------------------------------------------|-----------------------------------------------------------------------------------|
| <b>Two nodes</b>                         | <pre>ha-config modify controller mcc-2n<br/>ha-config modify chassis mcc-2n</pre> |
| <b>Four or eight nodes</b>               | <pre>ha-config modify controller mcc<br/>ha-config modify chassis mcc</pre>       |

## Reassign root aggregate disks

Reassign the root aggregate disks to the new controller module, using the sysids gathered earlier

### About this task

This task is performed in Maintenance mode.

The old system IDs were identified in [Gathering information before the upgrade](#).

The examples in this procedure use controllers with the following system IDs:

| Node     | Old system ID | New system ID |
|----------|---------------|---------------|
| node_B_1 | 4068741254    | 1574774970    |

### Steps

1. Cable all other connections to the new controller modules (FC-VI, storage, cluster interconnect, etc.).
2. Halt the system and boot to Maintenance mode from the LOADER prompt:

```
boot_ontap maint
```

3. Display the disks owned by node\_B\_1-A700:

```
disk show -a
```

The example output shows the system ID of the new controller module (1574774970). However, the root aggregate disks are still owned by the old system ID (4068741254). This example does not show drives owned by other nodes in the MetroCluster configuration.

```

*> disk show -a
Local System ID: 1574774970

 DISK OWNER POOL SERIAL NUMBER HOME
DR HOME

.....
...
rr18:9.126L44 node_B_1-A700(4068741254) Pool1 PZHYN0MD
node_B_1-A700(4068741254) node_B_1-A700(4068741254)
rr18:9.126L49 node_B_1-A700(4068741254) Pool1 PPG3J5HA
node_B_1-A700(4068741254) node_B_1-A700(4068741254)
rr18:8.126L21 node_B_1-A700(4068741254) Pool1 PZHTDSZD
node_B_1-A700(4068741254) node_B_1-A700(4068741254)
rr18:8.126L2 node_B_1-A700(4068741254) Pool0 SOM1J2CF
node_B_1-A700(4068741254) node_B_1-A700(4068741254)
rr18:8.126L3 node_B_1-A700(4068741254) Pool0 SOM0CQM5
node_B_1-A700(4068741254) node_B_1-A700(4068741254)
rr18:9.126L27 node_B_1-A700(4068741254) Pool0 SOM1PSDW
node_B_1-A700(4068741254) node_B_1-A700(4068741254)
...

```

4. Reassign the root aggregate disks on the drive shelves to the new controller:

```
disk reassign -s old-sysid -d new-sysid
```

The following example shows reassignment of drives:

```

*> disk reassign -s 4068741254 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? Jul 14 19:23:49
[localhost:config.bridge.extra.port:error]: Both FC ports of FC-to-SAS
bridge rtp-fc02-41-rr18:9.126L0 S/N [FB7500N107692] are attached to this
controller.
y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 4068741254.
Do you want to continue (y/n)? y

```

5. Check that all disks are reassigned as expected: `disk show`

```

*> disk show
Local System ID: 1574774970

 DISK OWNER POOL SERIAL NUMBER HOME
DR HOME

rr18:8.126L18 node_B_1-A900 (1574774970) Pool1 PZHYN0MD
node_B_1-A900 (1574774970) node_B_1-A900 (1574774970)
rr18:9.126L49 node_B_1-A900 (1574774970) Pool1 PPG3J5HA
node_B_1-A900 (1574774970) node_B_1-A900 (1574774970)
rr18:8.126L21 node_B_1-A900 (1574774970) Pool1 PZHTDSZD
node_B_1-A900 (1574774970) node_B_1-A900 (1574774970)
rr18:8.126L2 node_B_1-A900 (1574774970) Pool0 SOM1J2CF
node_B_1-A900 (1574774970) node_B_1-A900 (1574774970)
rr18:9.126L29 node_B_1-A900 (1574774970) Pool0 SOM0CQM5
node_B_1-A900 (1574774970) node_B_1-A900 (1574774970)
rr18:8.126L1 node_B_1-A900 (1574774970) Pool0 SOM1PSDW
node_B_1-A900 (1574774970) node_B_1-A900 (1574774970)
*>

```

6. Display the aggregate status: `aggr status`

```
*> aggr status
 Aggr State Status Options
aggr0_node_b_1-root online raid_dp, aggr root, nosnap=on,
 mirrored
mirror_resync_priority=high(fixed)
 fast zeroed
 64-bit
```

7. Repeat the above steps on the partner node (node\_B\_2-A900).

## Boot up the new controllers

You must reboot the controllers from the boot menu to update the controller flash image. Additional steps are required if encryption is configured.

### About this task

This task must be performed on all the new controllers.

### Steps

1. Halt the node: `halt`
2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address
```

```
setenv bootarg.kmip.init.netmask netmask
```

```
setenv bootarg.kmip.init.gateway gateway-address
```

```
setenv bootarg.kmip.init.interface interface-id
```

3. Display the boot menu: `boot_ontap menu`
4. If root encryption is used, issue the boot menu command for your key management configuration.

| If you are using...     | Select this boot menu option...                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option 10 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration |
| External key management | Option 11 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration |

5. If autoboot is enabled, interrupt autoboot by pressing control-C.
6. From the boot menu, run option (6).



Option 6 will reboot the node twice before completing.

Respond *y* to the system id change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...
Rebooting to load the restored env file...
```

7. Double-check that the partner-sysid is correct: `printenv partner-sysid`

If the partner-sysid is not correct, set it: `setenv partner-sysid partner-sysID`

8. If root encryption is used, issue the boot menu command again for your key management configuration.

| If you are using...     | Select this boot menu option...                                                                                     |
|-------------------------|---------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option 10 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration |
| External key management | Option 11 and follow the prompts to provide the required inputs to recover or restore the key-manager configuration |

You might need to issue the `recover_XXXXXXXX_keymanager` command at the boot menu prompt multiple times until the nodes completely boot.

9. Boot the nodes: `boot_ontap`

10. Wait for the replaced nodes to boot up.

If either node is in takeover mode, perform a giveback using the `storage failover giveback` command.

11. Verify that all ports are in a broadcast domain:

a. View the broadcast domains:

```
network port broadcast-domain show
```

b. Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

c. Add the physical port that will host the intercluster LIFs to the corresponding Broadcast domain.

d. Modify intercluster LIFs to use the new physical port as home port.

e. After the intercluster LIFs are up, check the cluster peer status and re-establish cluster peering as needed.

You may need to reconfigure cluster peering.

[Creating a cluster peer relationship](#)

f. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Creating a VLAN](#)

[Combining physical ports to create interface groups](#)

12. If encryption is used, restore the keys using the correct command for your key management configuration.

| If you are using...     | Use this command...                                                                                                                                                                                                                                          |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | <pre>security key-manager onboard sync</pre> <p>For more information, see <a href="#">Restoring onboard key management encryption keys</a>.</p>                                                                                                              |
| External key management | <pre>security key-manager external restore -vserver SVM -node node -key-server host_name IP_address:port -key-id key_id -key-tag key_tag node-name</pre> <p>For more information, see <a href="#">Restoring external key management encryption keys</a>.</p> |

## Verify LIF configuration

Verify that LIFs are hosted on appropriate node/ports prior to switchback. The following steps need to be performed

### About this task

This task is performed on site\_B, where the nodes have been booted up with root aggregates.

### Steps

1. Verify that LIFs are hosted on the appropriate node and ports prior to switchback.
  - a. Change to the advanced privilege level:

```
set -privilege advanced
```

- b. Override the port configuration to ensure proper LIF placement:

```
vserver config override -command "network interface modify" -vserver vserver_name -home-port active_port_after_upgrade -lif lif_name -home-node new_node_name"
```

When entering the `network interface modify` command within the `vserver config override` command, you cannot use the tab autocomplete feature. You can create the `network interface modify` using autocomplete and then enclose it in the `vserver config override` command.

- c. Return to the admin privilege level:

```
set -privilege admin
```

2. Revert the interfaces to their home node:

```
network interface revert * -vserver vservice-name
```

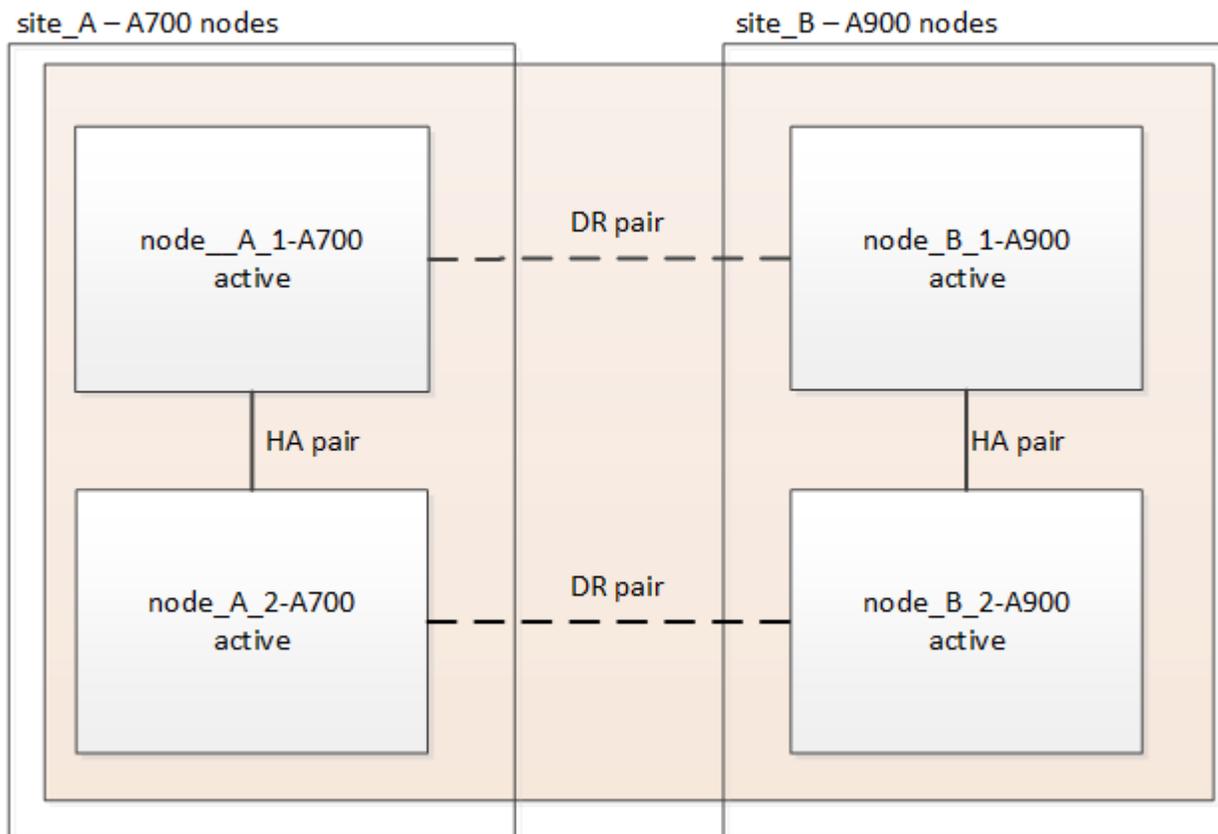
Perform this step on all SVMs as required.

## Switch back the MetroCluster configuration

After the new controllers have been configured, you switch back the MetroCluster configuration to return the configuration to normal operation.

### About this task

In this task, you will perform the switchback operation, returning the MetroCluster configuration to normal operation. The nodes on site\_A are still awaiting upgrade as shown in the following illustration. (This illustration also applies to upgrading a FAS9000 to a FAS9500 controller).



### Steps

1. Issue the `metrocluster node show` command on site\_B and check the output.
  - a. Verify that the new nodes are represented correctly.
  - b. Verify that the new nodes are in "Waiting for switchback state."
2. Switchback the cluster:

```
metrocluster switchback
```

### 3. Check the progress of the switchback operation:

```
metrocluster show
```

The switchback operation is still in progress when the output displays `waiting-for-switchback`:

```
cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
 Mode switchover
 AUSO Failure Domain -
Remote: cluster_A Configuration state configured
 Mode waiting-for-switchback
 AUSO Failure Domain -
```

The switchback operation is complete when the output displays `normal`:

```
cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
 Mode normal
 AUSO Failure Domain -
Remote: cluster_A Configuration state configured
 Mode normal
 AUSO Failure Domain -
```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command. This command is at the advanced privilege level.

## Check the health of the MetroCluster configuration

After upgrading the controller modules you must verify the health of the MetroCluster configuration.

### About this task

This task can be performed on any node in the MetroCluster configuration.

### Steps

#### 1. Verify the operation of the MetroCluster configuration:

- a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Perform a MetroCluster check:

```
metrocluster check run
```

- c. Display the results of the MetroCluster check:

```
metrocluster check show
```

After running the `metrocluster check run` and `metrocluster check show` commands, you might see an error similar to the following example:

```
Cluster_A:: node_A_1 (non-overridable veto): DR partner NVLog
mirroring is not online. Make sure that the links between the two
sites are healthy and properly configured.
```

This error occurs due to a controller mismatch during the upgrade process. You can safely ignore the error and proceed to upgrade the nodes on site\_A.

## Upgrade the nodes on site\_A

You must repeat the upgrade tasks on site\_A.

### Step

1. Repeat the steps to upgrade the nodes on site\_A, beginning with [Prepare for the upgrade](#).

As you perform the tasks, all example references to the sites and nodes are reversed. For example, when the example is given to switchover from site\_A, you will switchover from Site\_B.

## Send a custom AutoSupport message after maintenance

After completing the upgrade, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

### Step

1. To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.
  - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

## Restore Tiebreaker monitoring

If the MetroCluster configuration was previously configured for monitoring by the Tiebreaker software, you can restore the Tiebreaker connection.

1. Use the steps in: [Adding MetroCluster configurations](#) in the *MetroCluster Tiebreaker Installation and Configuration* section.

# Upgrade controllers in a four-node MetroCluster FC configuration using switchover and switchback with "system controller replace" commands (ONTAP 9.10.1 and later)

You can use this guided automated MetroCluster switchover operation to perform a non-disruptive controller upgrade on a four-node MetroCluster FC configuration. Other components (such as storage shelves or switches) cannot be upgraded as part of this procedure.

## Supported platform combinations

- For information on what platform upgrade combinations are supported, review the MetroCluster FC upgrade table in [Choose a controller upgrade procedure](#).

Refer to [Choose an upgrade or refresh method](#) for additional procedures.

## About this task

- You can use this procedure only for controller upgrade.

Other components in the configuration, such as storage shelves or switches, cannot be upgraded at the same time.

- This procedure applies to controller modules in a four-node MetroCluster FC configuration.
- The platforms must be running ONTAP 9.10.1 or later.

### [NetApp Hardware Universe](#)

- You can use this procedure to upgrade controllers in a four-node MetroCluster FC configuration using NSO based automated switchover and switchback. If you want to perform a controller upgrade using aggregate relocation (ARL), refer to [Use "system controller replace" commands to upgrade controller hardware running ONTAP 9.8 or later](#). It is recommended to use the NSO based automated procedure.
- If your MetroCluster sites are physically at two different locations, you should use the automated NSO controller upgrade procedure to upgrade the controllers at both sites in sequence.
- This automated NSO based controller upgrade procedure gives you the capability to initiate controller replacement to a MetroCluster disaster recovery (DR) site. You can only initiate a controller replacement at one site at a time.
- To initiate a controller replacement at site A, you need to run the controller replacement start command from site B. The operation guides you to replace controllers of both the nodes at site A only. To replace the controllers at site B, you need to run the controller replacement start command from site A. A message displays identifying the site at which the controllers are being replaced.

The following example names are used in this procedure:

- site\_A
  - Before upgrade:
    - node\_A\_1-old

- node\_A\_2-old
- After upgrade:
  - node\_A\_1-new
  - node\_A\_2-new
- site\_B
  - Before upgrade:
    - node\_B\_1-old
    - node\_B\_2-old
  - After upgrade:
    - node\_B\_1-new
    - node\_B\_2-new

## Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Prepare for the upgrade

To prepare for the controller upgrade, you need to perform system prechecks and collect the configuration information.

At any stage during the upgrade, you can run the `system controller replace show` or `system controller replace show-details` command from site A to check the status. If the commands return a blank output, wait for a few minutes and rerun the command.

### Steps

1. Run the following command from site A to replace the controllers at site B:

```
system controller replace start
```



- If you're repeating the procedure at one site, after already replacing the controllers at the other site, an error occurs due to a mismatch between the nodes at each site. This is the expected behavior when there are different platform models on both sites.

If only the mismatch error is returned, you can use the `-skip-metrocluster-check true` option with the `system controller replace start` command to skip the MetroCluster checks.

The automated operation executes the checks. If no issues are found, the operation pauses so you can manually collect the configuration related information.

The current source system and all compatible target systems are displayed. If you have replaced the source controller with a controller that has a different ONTAP version or a non-compatible platform, the automation operation halts and reports an error after the new nodes boot. To bring the cluster back to a healthy state, follow the manual recovery procedure.

The `system controller replace start` command might report the following precheck error:

```
Cluster-A::*>system controller replace show
Node Status Error-Action

Node-A-1 Failed MetroCluster check failed. Reason : MCC check
showed errors in component aggregates
```

Check if this error occurred because you have unmirrored aggregates or due to another aggregate issue. Verify that all mirrored aggregates are healthy and not degraded or mirror-degraded. If this error is due to unmirrored aggregates only, you can override this error by selecting the `-skip-metrocluster-check true` option on the `system controller replace start` command. If remote storage is accessible, the unmirrored aggregates come online after switchover. If the remote storage link fails, the unmirrored aggregates fail to come online.

2. Manually collect the configuration information by logging in at site B and following the commands listed in the console message under the `system controller replace show` or `system controller replace show-details` command.

## Gather information before the upgrade

Before upgrading, if the root volume is encrypted, you must gather the backup key and other information to boot the new controllers with the old encrypted root volumes.

### About this task

This task is performed on the existing MetroCluster FC configuration.

### Steps

1. Label the cables for the existing controllers, so you can easily identify the cables when setting up the new controllers.
2. Display the commands to capture the backup key and other information:

```
system controller replace show
```

Run the commands listed under the `show` command from the partner cluster.

### 3. Gather the system IDs of the nodes in the MetroCluster configuration:

```
metrocluster node show -fields node-systemid,dr-partner-systemid
```

During the upgrade procedure, you will replace these old system IDs with the system IDs of the new controller modules.

In this example for a four-node MetroCluster FC configuration, the following old system IDs are retrieved:

- `node_A_1-old: 4068741258`
- `node_A_2-old: 4068741260`
- `node_B_1-old: 4068741254`
- `node_B_2-old: 4068741256`

```
metrocluster-siteA::> metrocluster node show -fields node-systemid,ha-
partner-systemid,dr-partner-systemid,dr-auxiliary-systemid
dr-group-id cluster node node-systemid
ha-partner-systemid dr-partner-systemid dr-auxiliary-systemid

1 Cluster_A Node_A_1-old 4068741258
4068741260 4068741256
1 Cluster_A Node_A_2-old 4068741260
4068741258 4068741254
1 Cluster_B Node_B_1-old 4068741254
4068741256 4068741258
1 Cluster_B Node_B_2-old 4068741256
4068741254 4068741260
4 entries were displayed.
```

In this example for a two-node MetroCluster FC configuration, the following old system IDs are retrieved:

- `node_A_1: 4068741258`
- `node_B_1: 4068741254`

```
metrocluster node show -fields node-systemid,dr-partner-systemid
dr-group-id cluster node node-systemid dr-partner-systemid

1 Cluster_A Node_A_1-old 4068741258 4068741254
1 Cluster_B node_B_1-old - -
2 entries were displayed.
```

### 4. Gather port and LIF information for each old node.

You should gather the output of the following commands for each node:

- `network interface show -role cluster,node-mgmt`
- `network port show -node node-name -type physical`
- `network port vlan show -node node-name`
- `network port ifgrp show -node node_name -instance`
- `network port broadcast-domain show`
- `network port reachability show -detail`
- `network ipspace show`
- `volume show`
- `storage aggregate show`
- `system node run -node node-name sysconfig -a`

5. If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

You should gather the output of the following commands:

- `fc adapter show -instance`
- `fc interface show -instance`
- `iscsi interface show`
- `ucadmin show`

6. If the root volume is encrypted, collect and save the passphrase used for key-manager:

```
security key-manager backup show
```

7. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.

For additional information, see [Backing up onboard key management information manually](#).

a. If Onboard Key Manager is configured:

```
security key-manager onboard show-backup
```

You will need the passphrase later in the upgrade procedure.

b. If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
```

```
security key-manager key query
```

8. After you finish collecting the configuration information, resume the operation:

```
system controller replace resume
```

## Remove the existing configuration from the Tiebreaker or other monitoring software

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the Tiebreaker or other software prior to replacing the old controller.

### Steps

1. [Remove the existing MetroCluster configuration](#) from the Tiebreaker software.
2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

## Replace the old controllers and boot up the new controllers

After you gather information and resume the operation, the automation proceeds with the switchover operation.

### About this task

The automation operation initiates the switchover, `heal-aggregates`, and `heal root-aggregates` operations. After these operations complete, the operation pauses at **paused for user intervention** so you can rack and install the controllers, boot up the partner controllers, and reassign the root aggregate disks to the new controller module from flash backup using the `sysids` gathered earlier.

### Before you begin

Before initiating switchover, the automation operation pauses so you can manually verify that all LIFs are “up” at site B. If necessary, bring any LIFs that are “down” to “up” and resume the automation operation by using the `system controller replace resume` command.

### Prepare the network configuration of the old controllers

To ensure that the networking resumes cleanly on the new controllers, you must move LIFs to a common port and then remove the networking configuration of the old controllers.

### About this task

- This task must be performed on each of the old nodes.
- You will use the information gathered in [Prepare for the upgrade](#).

### Steps

1. Boot the old nodes and then log in to the nodes:

```
boot_ontap
```

2. Assign the home port of all data LIFs on the old controller to a common port that is the same on both the old and new controller modules.

- a. Display the LIFs:

```
network interface show
```

All data LIFs including SAN and NAS will be admin “up” and operationally “down” since those are up at switchover site (`cluster_A`).

- b. Review the output to find a common physical network port that is the same on both the old and new

controllers that is not used as a cluster port.

For example, “e0d” is a physical port on old controllers and is also present on new controllers. “e0d” is not used as a cluster port or otherwise on the new controllers.

For port usage for platform models, see the [NetApp Hardware Universe](#)

- c. Modify all data LIFS to use the common port as the home port:

```
network interface modify -vserver svm-name -lif data-lif -home-port port-id
```

In the following example, this is “e0d”.

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

3. Modify broadcast domains to remove VLAN and physical ports that need to be deleted:

```
broadcast-domain remove-ports -broadcast-domain broadcast-domain-name -ports
node-name:port-id
```

Repeat this step for all VLAN and physical ports.

4. Remove any VLAN ports using cluster ports as member ports and interface groups using cluster ports as member ports.

- a. Delete VLAN ports:

```
network port vlan delete -node node-name -vlan-name portid-vlandid
```

For example:

```
network port vlan delete -node node1 -vlan-name e1c-80
```

- b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node node-name -ifgrp interface-group-name
-port portid
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp ala -port e0d
```

- c. Remove VLAN and interface group ports from broadcast domain:

```
network port broadcast-domain remove-ports -ipSPACE ipSPACE -broadcast
-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..
```

- d. Modify interface group ports to use other physical ports as member as needed.:

```
ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id
```

5. Halt the nodes:

```
halt -inhibit-takeover true -node node-name
```

This step must be performed on both nodes.

## Set up the new controllers

You must rack and cable the new controllers.

### Steps

1. Plan out the positioning of the new controller modules and storage shelves as needed.

The rack space depends on the platform model of the controller modules, the switch types, and the number of storage shelves in your configuration.

2. Properly ground yourself.

3. Install the controller modules in the rack or cabinet.

[ONTAP Hardware Systems Documentation](#)

4. If the new controller modules did not come with FC-VI cards of their own and if FC-VI cards from old controllers are compatible on new controllers, swap FC-VI cards and install those in correct slots.

See the [NetApp Hardware Universe](#) for slot info for FC-VI cards.

5. Cable the controllers' power, serial console and management connections as described in the *MetroCluster Installation and Configuration Guides*.

Do not connect any other cables that were disconnected from old controllers at this time.

[ONTAP Hardware Systems Documentation](#)

6. Power up the new nodes and press Ctrl-C when prompted to display the LOADER prompt.

## Netboot the new controllers

After you install the new nodes, you need to netboot to ensure the new nodes are running the same version of ONTAP as the original nodes. The term netboot means you are booting from an ONTAP image stored on a remote server. When preparing for netboot, you must put a copy of the ONTAP 9 boot image onto a web server that the system can access.

This task is performed on each of the new controller modules.

### Steps

1. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.

2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `ontap-version_image.tgz` file on a web-accessible directory.

3. Go to the web-accessible directory and verify that the files you need are available.

| If the platform model is... | Then...                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FAS/AFF8000 series systems  | <p>Extract the contents of the ontap-version_image.tgzfile to the target directory: <code>tar -zxvf ontap-version_image.tgz</code></p> <p>NOTE: If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</p> <p>Your directory listing should contain a netboot folder with a kernel file:netboot/kernel</p> |
| All other systems           | <p>Your directory listing should contain a netboot folder with a kernel file: ontap-version_image.tgz</p> <p>You do not need to extract the ontap-version_image.tgz file.</p>                                                                                                                                                                        |

4. At the LOADER prompt, configure the netboot connection for a management LIF:

- If IP addressing is DHCP, configure the automatic connection:

```
ifconfig e0M -auto
```

- If IP addressing is static, configure the manual connection:

```
ifconfig e0M -addr=ip_addr -mask=netmask -gw=gateway
```

5. Perform the netboot.

- If the platform is an 80xx series system, use this command:

```
netboot http://web_server_ip/path_to_web-accessible_directory/netboot/kernel
```

- If the platform is any other system, use the following command:

```
netboot http://web_server_ip/path_to_web-accessible_directory/ontap-version_image.tgz
```

6. From the boot menu, select option **(7) Install new software first** to download and install the new software image to the boot device.

Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". It applies to nondisruptive upgrades of software, not to upgrades of controllers.

7. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL of the image file: `http://web_server_ip/path_to_web-accessible_directory/ontap-version_image.tgz`

```
Enter username/password if applicable, or press Enter to continue.
```

8. Enter `n` to skip the backup recovery when you see a prompt similar to the following:

```
Do you want to restore the backup configuration now? {y|n} n
```

9. Reboot by entering `y` when you see a prompt similar to the following:

```
The node must be rebooted to start using the newly installed software.
Do you want to reboot now? {y|n} y
```



You must reboot the node in order to use the newly installed software.

### Clear the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing configuration.

#### Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

## Restore the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site's usage.

### Steps

1. In Maintenance mode configure the settings for any HBAs in the system:

- a. Check the current settings of the ports: `ucadmin show`
- b. Update the port settings as needed.

| If you have this type of HBA and desired mode... | Use this command...                                         |
|--------------------------------------------------|-------------------------------------------------------------|
| CNA FC                                           | <code>ucadmin modify -m fc -t initiator adapter-name</code> |
| CNA Ethernet                                     | <code>ucadmin modify -mode cna adapter-name</code>          |
| FC target                                        | <code>fcadmin config -t target adapter-name</code>          |
| FC initiator                                     | <code>fcadmin config -t initiator adapter-name</code>       |

2. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

3. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

```
boot_ontap maint
```

4. Verify the changes you made:

| If you have this type of HBA... | Use this command...       |
|---------------------------------|---------------------------|
| CNA                             | <code>ucadmin show</code> |
| FC                              | <code>fcadmin show</code> |

## Reassign root aggregate disks

Reassign the root aggregate disks to the new controller module, using the `sysids` gathered earlier

### About this task

This task is performed in Maintenance mode.

The old system IDs were identified in [Gather information before the upgrade](#).

The examples in this procedure use controllers with the following system IDs:

| Node     | Old system ID | New system ID |
|----------|---------------|---------------|
| node_B_1 | 4068741254    | 1574774970    |

### Steps

1. Cable all other connections to the new controller modules (FC-VI, storage, cluster interconnect, etc.).
2. Halt the system and boot to Maintenance mode from the LOADER prompt:

```
boot_ontap maint
```

3. Display the disks owned by node\_B\_1-old:

```
disk show -a
```

The command output shows the system ID of the new controller module (1574774970). However, the root aggregate disks are still owned by the old system ID (4068741254). This example does not show drives owned by other nodes in the MetroCluster configuration.

```
*> disk show -a
Local System ID: 1574774970

 DISK OWNER POOL SERIAL NUMBER HOME
DR HOME

.....
...
rr18:9.126L44 node_B_1-old(4068741254) Pool1 PZHYN0MD
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:9.126L49 node_B_1-old(4068741254) Pool1 PPG3J5HA
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L21 node_B_1-old(4068741254) Pool1 PZHTDSZD
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L2 node_B_1-old(4068741254) Pool10 SOM1J2CF
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:8.126L3 node_B_1-old(4068741254) Pool10 SOM0CQM5
node_B_1-old(4068741254) node_B_1-old(4068741254)
rr18:9.126L27 node_B_1-old(4068741254) Pool10 SOM1PSDW
node_B_1-old(4068741254) node_B_1-old(4068741254)
...
```

4. Reassign the root aggregate disks on the drive shelves to the new controller:

```
disk reassign -s old-sysid -d new-sysid
```

The following example shows reassignment of drives:

```

*> disk reassign -s 4068741254 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? Jul 14 19:23:49
[localhost:config.bridge.extra.port:error]: Both FC ports of FC-to-SAS
bridge rtp-fc02-41-rr18:9.126L0 S/N [FB7500N107692] are attached to this
controller.
y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 4068741254.
Do you want to continue (y/n)? y

```

##### 5. Check that all disks are reassigned as expected:

```
disk show
```

```

*> disk show
Local System ID: 1574774970

 DISK OWNER POOL SERIAL NUMBER HOME
DR HOME

rr18:8.126L18 node_B_1-new(1574774970) Pool1 PZHYN0MD
node_B_1-new(1574774970) node_B_1-new(1574774970)
rr18:9.126L49 node_B_1-new(1574774970) Pool1 PPG3J5HA
node_B_1-new(1574774970) node_B_1-new(1574774970)
rr18:8.126L21 node_B_1-new(1574774970) Pool1 PZHTDSZD
node_B_1-new(1574774970) node_B_1-new(1574774970)
rr18:8.126L2 node_B_1-new(1574774970) Pool0 SOM1J2CF
node_B_1-new(1574774970) node_B_1-new(1574774970)
rr18:9.126L29 node_B_1-new(1574774970) Pool0 SOM0CQM5
node_B_1-new(1574774970) node_B_1-new(1574774970)
rr18:8.126L1 node_B_1-new(1574774970) Pool0 SOM1PSDW
node_B_1-new(1574774970) node_B_1-new(1574774970)
*>

```

## 6. Display the aggregate status:

```
aggr status
```

```
*> aggr status
 Aggr State Status Options
aggr0_node_b_1-root online raid_dp, aggr root, nosnap=on,
 mirrored
mirror_resync_priority=high(fixed)
 fast zeroed
 64-bit
```

## 7. Repeat the above steps on the partner node (node\_B\_2-new).

### Boot up the new controllers

You must reboot the controllers from the boot menu to update the controller flash image. Additional steps are required if encryption is configured.

You can reconfigure VLANs and interface groups. If required, manually modify the ports for the cluster LIFs and broadcast domain details before resuming the operation by using the `system controller replace resume` command.

### About this task

This task must be performed on all the new controllers.

### Steps

#### 1. Halt the node:

```
halt
```

#### 2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address
```

```
setenv bootarg.kmip.init.netmask netmask
```

```
setenv bootarg.kmip.init.gateway gateway-address
```

```
setenv bootarg.kmip.init.interface interface-id
```

#### 3. Display the boot menu:

```
boot_ontap menu
```

#### 4. If root encryption is used, select the boot menu option for your key management configuration.

| If you are using... | Select this boot menu option... |
|---------------------|---------------------------------|
|---------------------|---------------------------------|

|                         |                                                                                                                            |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option “10”<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option “11”<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |

- If autoboot is enabled, interrupt autoboot by pressing Ctrl-C.
- From the boot menu, run option “6”.



Option “6” will reboot the node twice before completing.

Respond “y” to the system id change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...

Rebooting to load the restored env file...
```

- Double-check that the partner-sysid is correct:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid partner-sysID
```

- If root encryption is used, select the boot menu option again for your key management configuration.

| If you are using...     | Select this boot menu option...                                                                                            |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option “10”<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option “11”<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |

Depending on the key manager setting, perform the recovery procedure by selecting option “10” or option “11”, followed by option “6” at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option “1” (normal boot).

- Boot the nodes:

```
boot_ontap
```

10. Wait for the replaced nodes to boot up.

If either node is in takeover mode, perform a giveback using the `storage failover giveback` command.

11. Verify that all ports are in a broadcast domain:

a. View the broadcast domains:

```
network port broadcast-domain show
```

b. Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

c. Add the physical port that will host the intercluster LIFs to the corresponding broadcast domain.

d. Modify intercluster LIFs to use the new physical port as home port.

e. After the intercluster LIFs are up, check the cluster peer status and re-establish cluster peering as needed.

You may need to reconfigure cluster peering.

[Create a cluster peer relationship](#)

f. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

[Create a VLAN](#)

[Combine physical ports to create interface groups](#)

g. Verify that the partner cluster is reachable and that the configuration is successfully resynchronized on the partner cluster:

```
metrocluster switchback -simulate true
```

12. If encryption is used, restore the keys using the correct command for your key management configuration.

| If you are using...    | Use this command...                                                                                                                           |
|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------|
| Onboard key management | <pre>security key-manager onboard sync</pre> <p>For more information, see <a href="#">Restore onboard key management encryption keys</a>.</p> |

## External key management

```
security key-manager external restore
-vserver SVM -node node -key-server
host_name|IP_address:port -key-id
key_id -key-tag key_tag node-name
```

For more information, see [Restore external key management encryption keys](#).

13. Before you resume the operation, verify that the MetroCluster is configured correctly. Check the node status:

```
metrocluster node show
```

Verify that the new nodes (site\_B) are in **Waiting for switchback state** from site\_A.

14. Resume the operation:

```
system controller replace resume
```

## Complete the upgrade

The automation operation runs verification system checks and then pauses so you can verify the network reachability. After verification, the resource regain phase is initiated and the automation operation executes switchback at site A and pauses at the post upgrade checks. After you resume the automation operation, it performs the post upgrade checks and if no errors are detected, marks the upgrade as complete.

### Steps

1. Verify the network reachability by following the console message.
2. After you complete the verification, resume the operation:

```
system controller replace resume
```

3. The automation operation performs switchback at site A and the post upgrade checks. When the operation pauses, manually check the SAN LIF status and verify the network configuration by following the console message.
4. After you complete the verification, resume the operation:

```
system controller replace resume
```

5. Check the post upgrade checks status:

```
system controller replace show
```

If the post upgrade checks did not report any errors, the upgrade is complete.

6. After you complete the controller upgrade, log in at site B and verify that the replaced controllers are configured correctly.

## Upgrade the nodes on cluster\_A

You must repeat the upgrade tasks to upgrade the nodes on cluster\_A at site A.

### Steps

1. Repeat the steps to upgrade the nodes on cluster\_A, beginning with [prepare for the upgrade](#).

When you repeat the procedure, all example references to the clusters and nodes are reversed.

### Restoring Tiebreaker monitoring

If the MetroCluster configuration was previously configured for monitoring by the Tiebreaker software, you can restore the Tiebreaker connection.

1. Use the steps in [Add MetroCluster configurations](#).

## Refreshing a four-node MetroCluster FC configuration

You can upgrade the controllers and storage in a four-node MetroCluster configuration by expanding the configuration to become an eight-node configuration and then removing the old disaster recovery (DR) group.

### About this task

References to "old nodes" mean the nodes that you intend to replace.

- You can only refresh specific platform models using this procedure in a MetroCluster FC configuration.
  - For information on what platform upgrade combinations are supported review the MetroCluster FC refresh table in [Choosing a system refresh method](#).

### Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

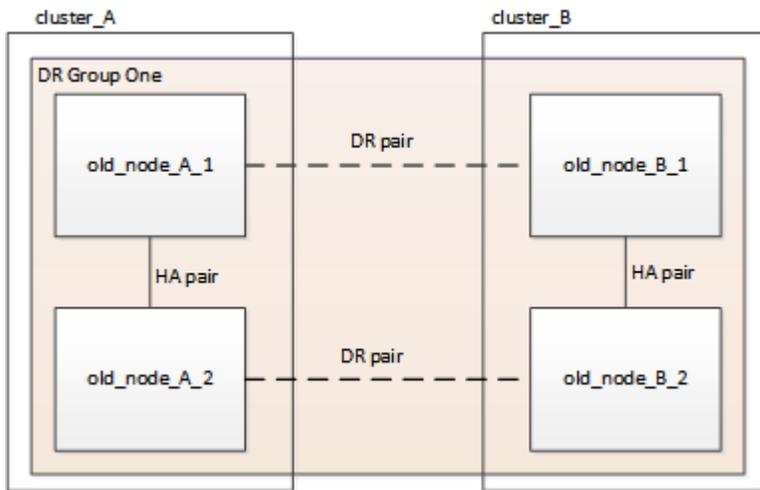
### Perform the refresh procedure

Use the following steps to refresh the MetroCluster FC configuration.

### Steps

1. Gather information from the old nodes.

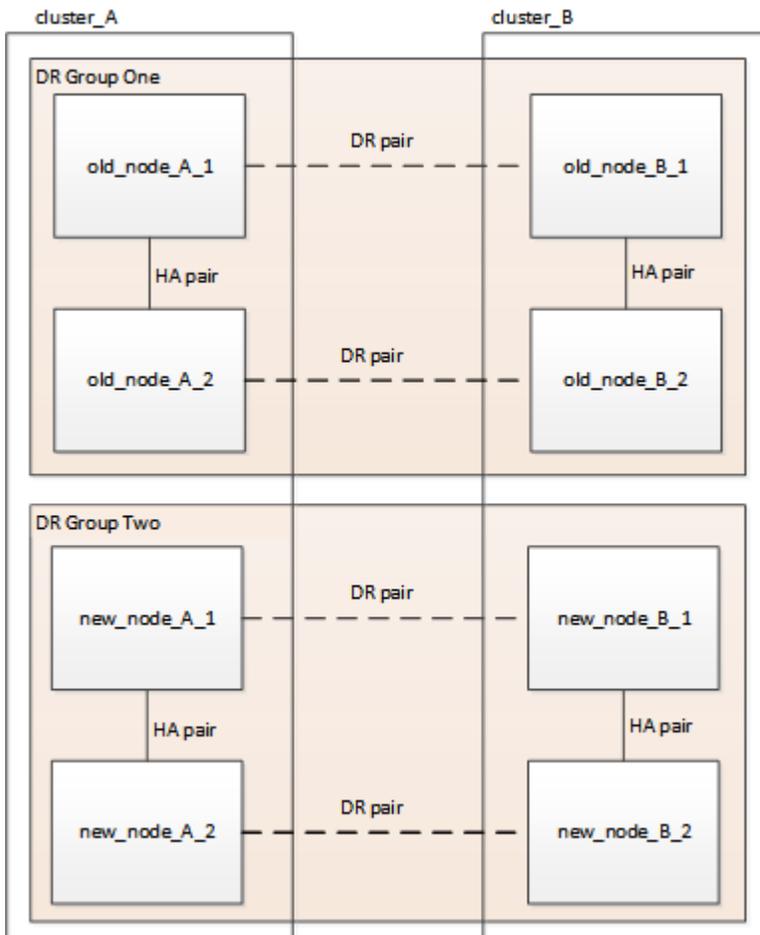
At this stage, the four-node configuration appears as shown in the following image:



2. Perform all of the steps in the four-node expansion procedure for your MetroCluster type.

[Expanding a four-node MetroCluster FC configuration to an eight-node configuration](#)

When the expansion procedure is complete, the configuration appears as shown in the following image:



3. Move the CRS volumes.

Perform the steps in [Move a metadata volume in MetroCluster configurations](#).

4. Move the data from the old nodes to new nodes using the following procedures:
  - a. Perform all the steps in [Create an aggregate and moving volumes to the new nodes](#).

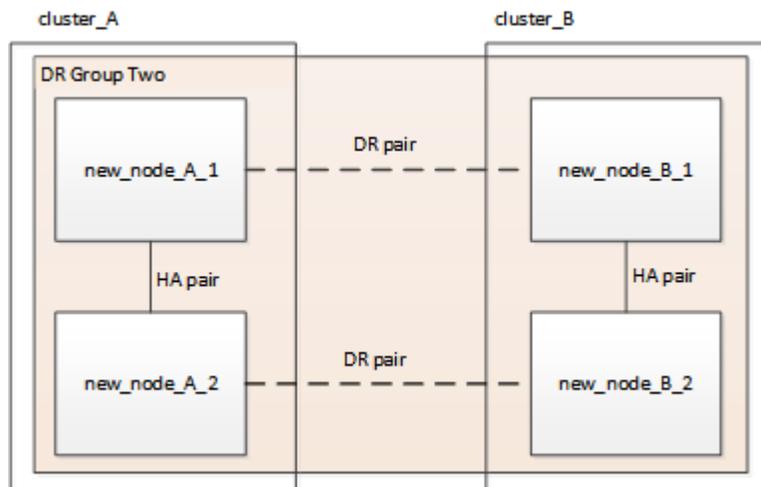


You might choose to mirror the aggregate when or after it is created.

- b. Perform all the steps in [Move non-SAN data LIFs and cluster management LIFs to the new nodes](#).
  - c. Perform all the steps in [Delete SAN LIFs no longer required from the original nodes](#).
5. Follow the steps in the procedure for removing the old DR group.

### Removing a Disaster Recovery group

After you have removed the old DR group (DR group one), the configuration appears as shown in the following image:



## Refresh a four-node or an eight-node MetroCluster IP configuration (ONTAP 9.8 and later)

You can use this procedure to upgrade controllers and storage in four-node or eight-node configurations.

Beginning with ONTAP 9.13.1, you can upgrade the controllers and storage in an eight-node MetroCluster IP configuration by expanding the configuration to become a temporary twelve-node configuration and then removing the old disaster recovery (DR) groups.

Beginning with ONTAP 9.8, you can upgrade the controllers and storage in a four-node MetroCluster IP configuration by expanding the configuration to become a temporary eight-node configuration and then removing the old DR group.

### Important information if you are adding an older platform model

The following guidance is for an uncommon scenario where you need to add an older platform model (platforms released before ONTAP 9.15.1) to an existing MetroCluster configuration that contains a newer platform model (platforms released in ONTAP 9.15.1 or later).

If your existing MetroCluster configuration contains a platform that uses **shared cluster/HA ports** (platforms released in ONTAP 9.15.1 or later), you cannot add a platform that uses **shared MetroCluster/HA ports**

(platforms released before ONTAP 9.15.1) without upgrading all nodes in the configuration to ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later.



Adding an older platform model that uses **shared/MetroCluster HA ports** to a MetroCluster containing a newer platform model that uses **shared cluster/HA ports** is an uncommon scenario and most combinations are not affected.

Use the following table to verify whether your combination is affected. If your existing platform is listed in the first column, and the platform you want to add to the configuration is listed in the second column, all nodes in the configuration must be running ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later to add the new DR group.

| If your existing MetroCluster contains..                                                                                                                                                                                                                                 |                                                                                                                                                 | And the platform you're adding is...                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                        | Then...                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| An AFF system using <b>shared cluster/HA ports</b> : <ul style="list-style-type: none"> <li>• AFF A20</li> <li>• AFF A30</li> <li>• AFF C30</li> <li>• AFF A50</li> <li>• AFF C60</li> <li>• AFF C80</li> <li>• AFF A70</li> <li>• AFF A90</li> <li>• AFF A1K</li> </ul> | A FAS system using <b>shared cluster/HA ports</b> : <ul style="list-style-type: none"> <li>• FAS50</li> <li>• FAS70</li> <li>• FAS90</li> </ul> | An AFF system using <b>shared MetroCluster/HA ports</b> : <ul style="list-style-type: none"> <li>• AFF A150, ASA A150</li> <li>• AFF A220</li> <li>• AFF C250, ASA C250</li> <li>• AFF A250, ASA A250</li> <li>• AFF A300</li> <li>• AFF A320</li> <li>• AFF C400, ASA C400</li> <li>• AFF A400, ASA A400</li> <li>• AFF A700</li> <li>• AFF C800, ASA C800</li> <li>• AFF A800, ASA A800</li> <li>• AFF A900, ASA A900</li> </ul> | A FAS system using <b>shared MetroCluster/HA ports</b> : <ul style="list-style-type: none"> <li>• FAS2750</li> <li>• FAS500f</li> <li>• FAS8200</li> <li>• FAS8300</li> <li>• FAS8700</li> <li>• FAS9000</li> <li>• FAS9500</li> </ul> | Before you add the new platform to your existing MetroCluster configuration, upgrade all nodes in the existing and new configuration to ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later. |

#### About this task

- If you have an eight-node configuration, your system must be running ONTAP 9.13.1 or later.
- If you have a four-node configuration, your system must be running ONTAP 9.8 or later.
- If you are also upgrading the IP switches, you must upgrade them before performing this refresh procedure.
- This procedure describes the steps required to refresh one four-node DR group. If you have an eight-node configuration (two DR groups) you can refresh one or both DR groups.

Refresh DR groups one at a time.

- \* References to "old nodes" mean the nodes that you intend to replace.
- \* For eight-node configurations, the source and target eight-node MetroCluster platform combination must be supported.



If you refresh both DR groups, the platform combination might not be supported after you refresh the first DR group. You must refresh both DR groups to achieve a supported eight-node configuration.

- You can only refresh specific platform models using this procedure in a MetroCluster IP configuration.
  - For information on which platform upgrade combinations are supported, review the MetroCluster IP refresh table in [Choosing a system refresh method](#).
- The lower limits of the source and target platforms apply. If you transition to a higher platform, the limits of the new platform applies only after the tech refresh of all DR groups completes.
- If you perform a tech refresh to a platform with lower limits than the source platform, you must adjust and reduce the limits to be at, or below, the target platform limits before performing this procedure.

## Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Perform the refresh procedure

Use the following steps to refresh the MetroCluster IP configuration.

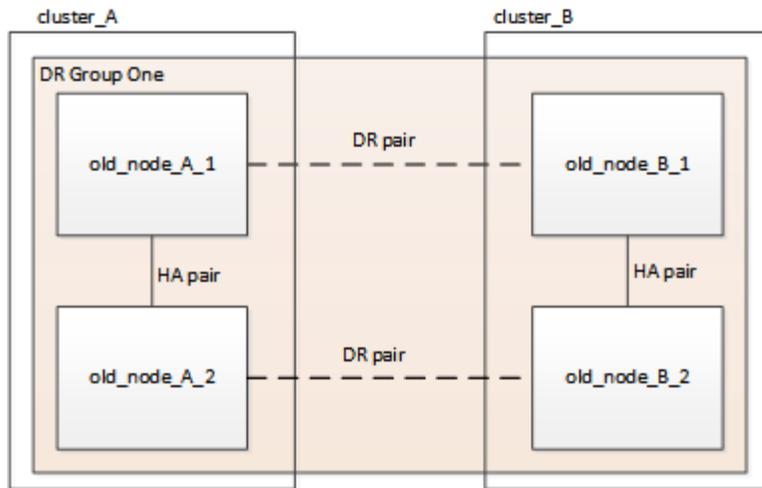
### Steps

1. Verify that you have a default broadcast domain created on the old nodes.

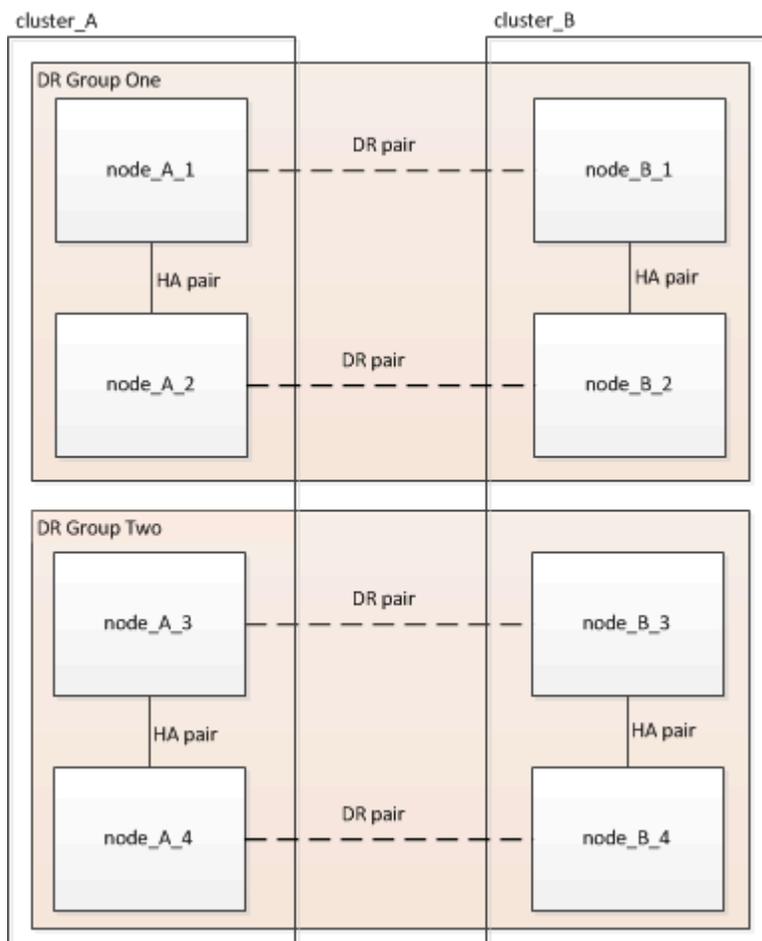
When you add new nodes to an existing cluster without a default broadcast domain, node management LIFs are created for the new nodes using universal unique identifiers (UUIDs) instead of the expected names. For more information, see the Knowledge Base article [Node management LIFs on newly-added nodes generated with UUID names](#).

2. Gather information from the old nodes.

At this stage, the four-node configuration appears as shown in the following image:



The eight-node configuration appears as shown in the following image:



3. To prevent automatic support case generation, send an AutoSupport message to indicate the upgrade is underway.

a. Issue the following command:

```
system node autosupport invoke -node * -type all -message "MAINT=10h
Upgrading old-model to new-model"
```

The following example specifies a 10 hour maintenance window. Allow additional time depending on

your plan.

If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

b. Repeat the command on the partner cluster.

4. If end-to-end encryption is enabled, follow the steps to [Disable end-to-end encryption](#).
5. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

| If you are using...      | Use this procedure...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tiebreaker               | <ol style="list-style-type: none"><li>1. Use the Tiebreaker CLI <code>monitor remove</code> command to remove the MetroCluster configuration.<br/><br/>In the following example, "cluster_A" is removed from the software:<br/><div data-bbox="889 827 1485 1167" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"><pre>NetApp MetroCluster Tiebreaker :&gt; monitor remove -monitor -name cluster_A Successfully removed monitor from NetApp MetroCluster Tiebreaker software.</pre></div></li><li>2. Confirm that the MetroCluster configuration is removed correctly by using the Tiebreaker CLI <code>monitor show -status</code> command.<br/><div data-bbox="889 1339 1485 1476" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"><pre>NetApp MetroCluster Tiebreaker :&gt; monitor show -status</pre></div></li></ol> |
| Mediator                 | Issue the following command from the ONTAP prompt:<br><br><pre>metrocluster configuration-settings mediator remove</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Third-party applications | Refer to the product documentation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

6. Perform all of the steps in [Expanding a MetroCluster IP configuration](#) to add the new nodes and storage to the configuration.

When the expansion procedure is complete, the temporary configuration appears as shown in the following images:

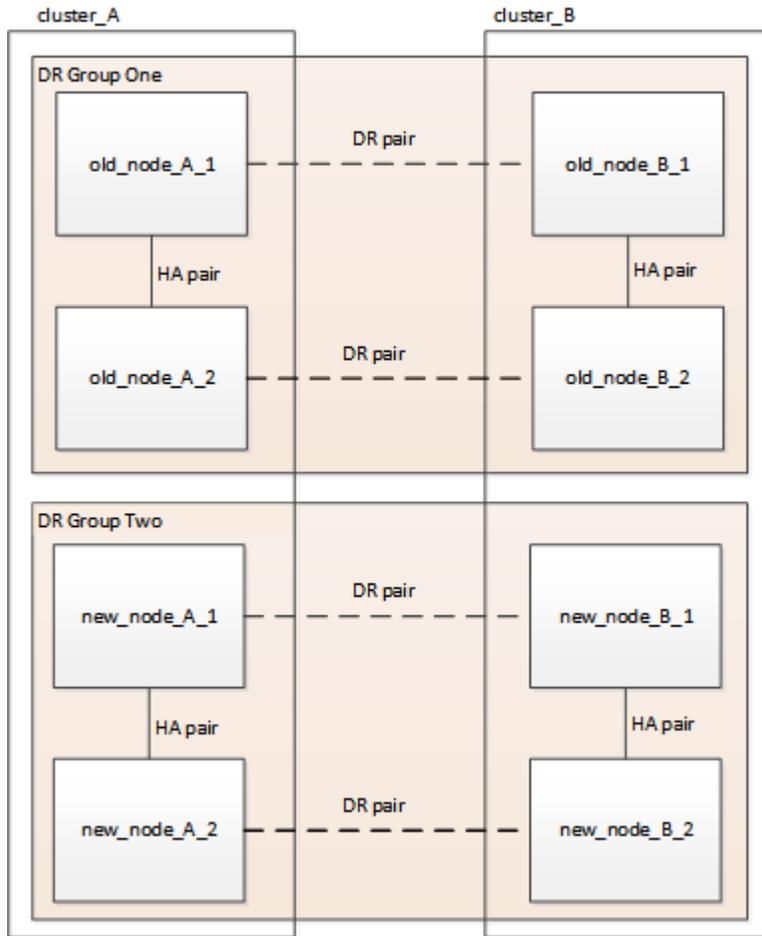
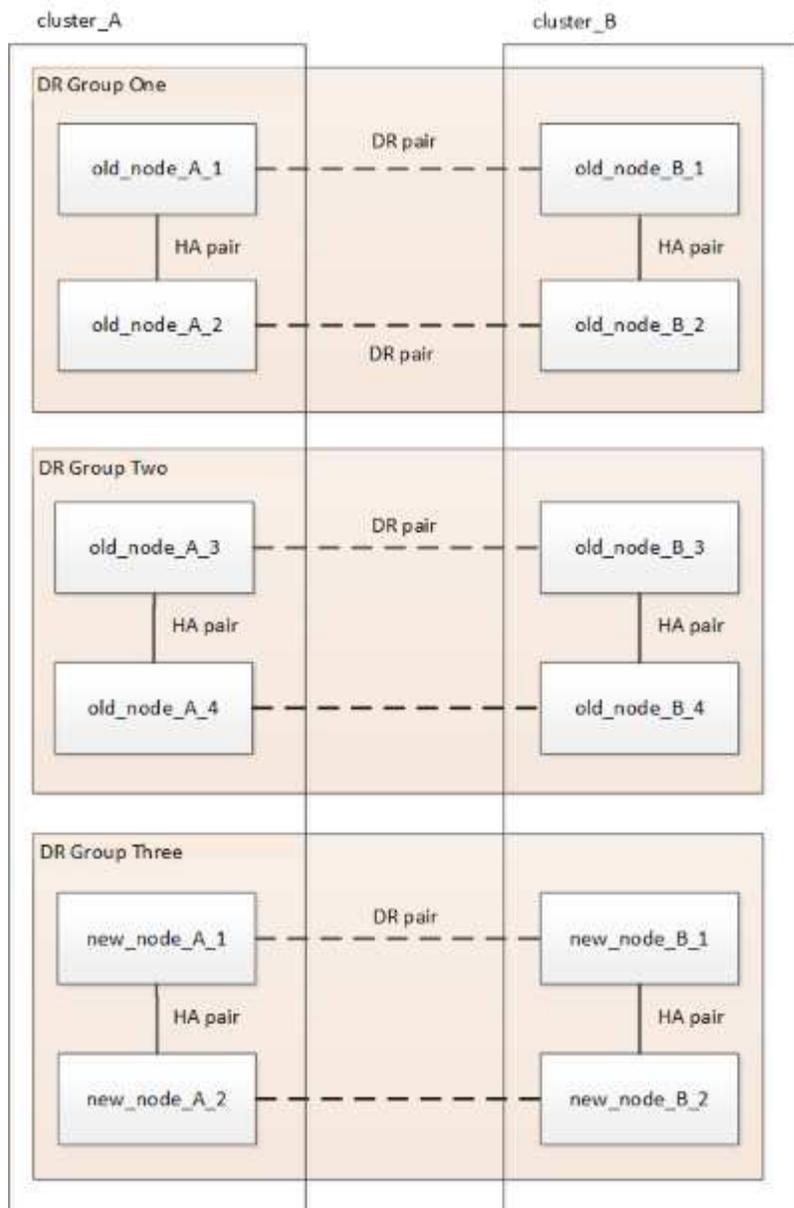


Figure 1. Temporary eight-node configuration



**Figure 2. Temporary twelve-node configuration**

7. Confirm that takeover is possible and the nodes are connected by running the following command on both clusters:

```
storage failover show
```

```
cluster_A::> storage failover show
```

| Node      | Partner   | Takeover Possible | State Description      |
|-----------|-----------|-------------------|------------------------|
| Node_FC_1 | Node_FC_2 | true              | Connected to Node_FC_2 |
| Node_FC_2 | Node_FC_1 | true              | Connected to Node_FC_1 |
| Node_IP_1 | Node_IP_2 | true              | Connected to Node_IP_2 |
| Node_IP_2 | Node_IP_1 | true              | Connected to Node_IP_1 |

8. Move the CRS volumes.

Perform the steps in [Moving a metadata volume in MetroCluster configurations](#).

9. Move the data from the old nodes to the new nodes by using the following procedures:

- a. Perform all the steps in [Create an aggregate and move volumes to the new nodes](#).



You might choose to mirror the aggregate when or after it is created.

- b. Perform all the steps in [Move non-SAN data LIFs and cluster-management LIFs to the new nodes](#).

10. Modify the IP address for the cluster peer of the transitioned nodes for each cluster:

- a. Identify the cluster\_A peer by using the `cluster peer show` command:

```
cluster_A::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability
Authentication

cluster_B 1-80-000011 Unavailable absent
```

- i. Modify the cluster\_A peer IP address:

```
cluster peer modify -cluster cluster_A -peer-addr node_A_3_IP -address
-family ipv4
```

- b. Identify the cluster\_B peer by using the `cluster peer show` command:

```
cluster_B::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability
Authentication

cluster_A 1-80-000011 Unavailable absent
```

- i. Modify the cluster\_B peer IP address:

```
cluster peer modify -cluster cluster_B -peer-addr node_B_3_IP -address
-family ipv4
```

- c. Verify that the cluster peer IP address is updated for each cluster:

- i. Verify that the IP address is updated for each cluster by using the `cluster peer show -instance` command.

The Remote Intercluster Addresses field in the following examples displays the updated IP address.

### Example for cluster\_A:

```
cluster_A::> cluster peer show -instance

Peer Cluster Name: cluster_B
 Remote Intercluster Addresses: 172.21.178.204,
172.21.178.212
 Availability of the Remote Cluster: Available
 Remote Cluster Name: cluster_B
 Active IP Addresses: 172.21.178.212,
172.21.178.204
 Cluster Serial Number: 1-80-000011
 Remote Cluster Nodes: node_B_3-IP,
node_B_4-IP
 Remote Cluster Health: true
 Unreachable Local Nodes: -
 Address Family of Relationship: ipv4
 Authentication Status Administrative: use-authentication
 Authentication Status Operational: ok
 Last Update Time: 4/20/2023 18:23:53
 IPspace for the Relationship: Default
Proposed Setting for Encryption of Inter-Cluster Communication: -
Encryption Protocol For Inter-Cluster Communication: tls-psk
 Algorithm By Which the PSK Was Derived: jpake

cluster_A::>
```

### Example for cluster\_B

```

cluster_B::> cluster peer show -instance

Peer Cluster Name: cluster_A
Remote Intercluster Addresses: 172.21.178.188,
172.21.178.196 <<<<<<< Should reflect the modified address
Availability of the Remote Cluster: Available
Remote Cluster Name: cluster_A
Active IP Addresses: 172.21.178.196,
172.21.178.188

Cluster Serial Number: 1-80-000011
Remote Cluster Nodes: node_A_3-IP,
node_A_4-IP

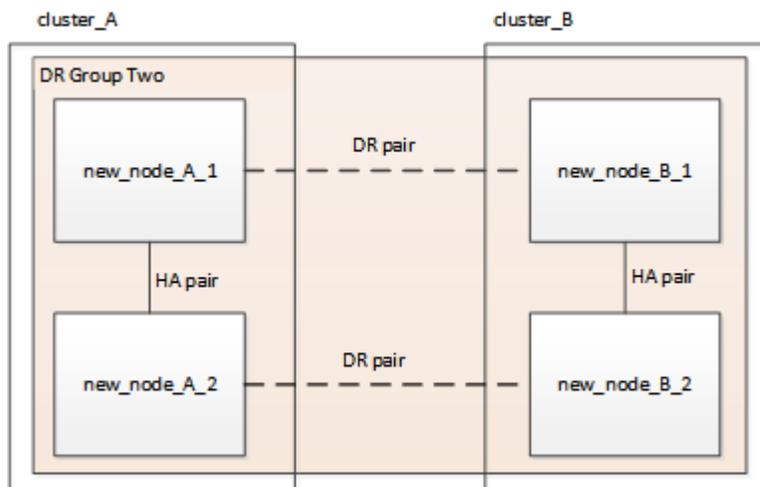
Remote Cluster Health: true
Unreachable Local Nodes: -
Address Family of Relationship: ipv4
Authentication Status Administrative: use-authentication
Authentication Status Operational: ok
Last Update Time: 4/20/2023 18:23:53
IPspace for the Relationship: Default
Proposed Setting for Encryption of Inter-Cluster Communication: -
Encryption Protocol For Inter-Cluster Communication: tls-psk
Algorithm By Which the PSK Was Derived: jpake

cluster_B::>

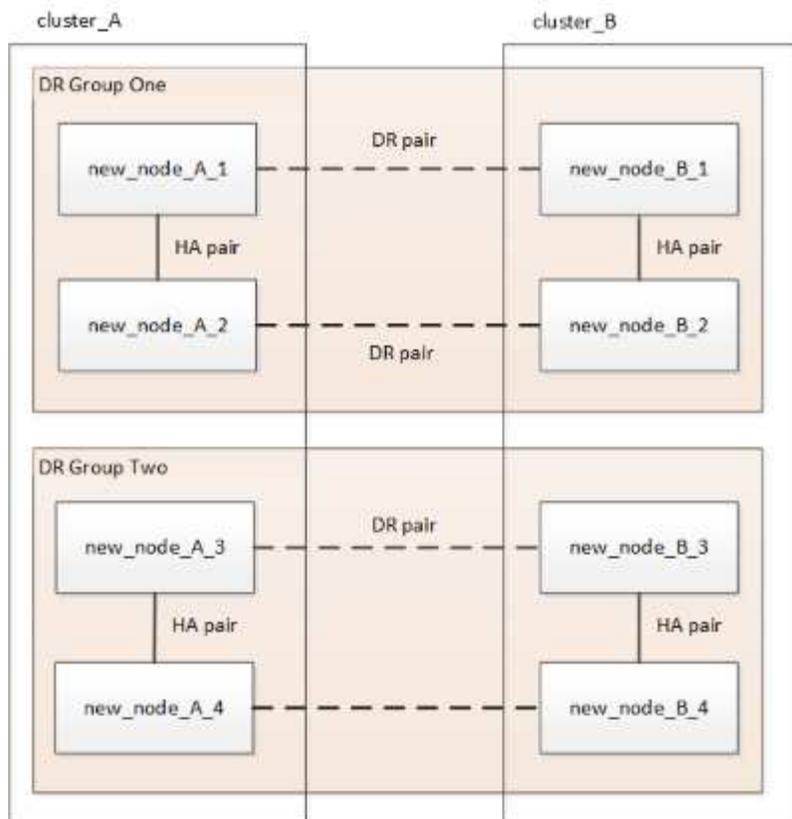
```

11. Follow the steps in [Removing a Disaster Recovery group](#) to remove the old DR group.
12. If you need to refresh both DR groups in an eight-node configuration, repeat the entire procedure for each DR group.

After you have removed the old DR group, the configuration appears as shown in the following images:



**Figure 3. Four-node configuration**



**Figure 4. Eight-node configuration**

13. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.
  - a. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- b. Confirm that all expected nodes are shown:

```
metrocluster node show
```

- c. Issue the following command:

```
metrocluster check run
```

- d. Display the results of the MetroCluster check:

```
metrocluster check show
```

14. If you disabled end-to-end encryption before adding the new nodes, you can re-enable it by following the steps in [Enable end-to-end encryption](#).
15. Restore monitoring if necessary, using the procedure for your configuration.

| If you are using... | Use this procedure                                                                                                        |
|---------------------|---------------------------------------------------------------------------------------------------------------------------|
| Tiebreaker          | <a href="#">Adding MetroCluster configurations</a> in the <i>MetroCluster Tiebreaker Installation and Configuration</i> . |

|                          |                                                                                                                                              |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| Mediator                 | <a href="#">Configure ONTAP Mediator from a MetroCluster IP configuration</a> in the <i>MetroCluster IP Installation and Configuration</i> . |
| Third-party applications | Refer to the product documentation.                                                                                                          |

16. To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.

a. Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

b. Repeat the command on the partner cluster.

## Expand a two-node MetroCluster FC configuration to a four-node configuration

### Expanding a two-node MetroCluster FC configuration to a four-node configuration

Expanding a two-node MetroCluster FC configuration to a four-node MetroCluster FC configuration involves adding a controller to each cluster to form an HA pair at each MetroCluster site, and then refreshing the MetroCluster FC configuration.

#### Before you begin

- The nodes must be running ONTAP 9 or later in a MetroCluster FC configuration.

This procedure is not supported on earlier versions of ONTAP or in MetroCluster IP configurations.

- If the platforms in your two-node configuration are not supported in ONTAP 9.2 and you plan to upgrade to platforms supported in ONTAP 9.2 *and* expand to a four-node cluster, you must upgrade the platforms in the two-node configuration *before* expanding the MetroCluster FC configuration.
- The existing MetroCluster FC configuration must be healthy.
- The equipment you are adding must be supported and meet all of the requirements described in the following procedures:

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

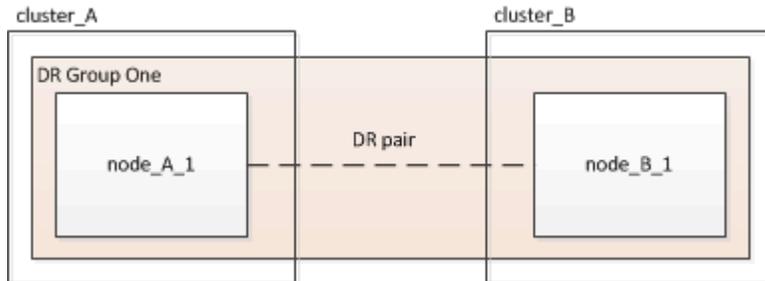
- You must have available FC switch ports to accommodate the new controllers and any new bridges.
- Verify that you have a default broadcast domain created on the old nodes.

When you add new nodes to an existing cluster without a default broadcast domain, node management LIFs are created for the new nodes using universal unique identifiers (UUIDs) instead of the expected names. For more information, see the Knowledge Base article [Node management LIFs on newly-added nodes generated with UUID names](#).

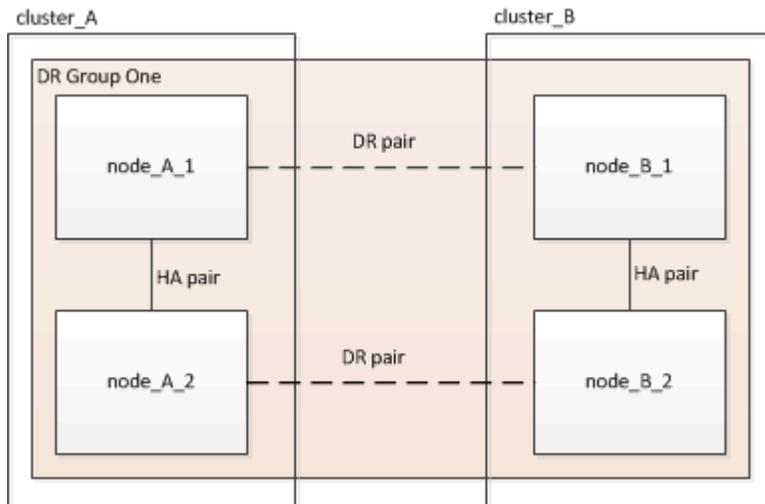
- You need the admin password and access to an FTP or SCP server.

## About this task

- This procedure applies only to MetroCluster FC configurations.
- This procedure is disruptive and takes approximately four hours to complete.
- Before performing this procedure, the MetroCluster FC configuration consists of two single-node clusters:



After completing this procedure, the MetroCluster FC configuration consists of two HA pairs, one at each site:



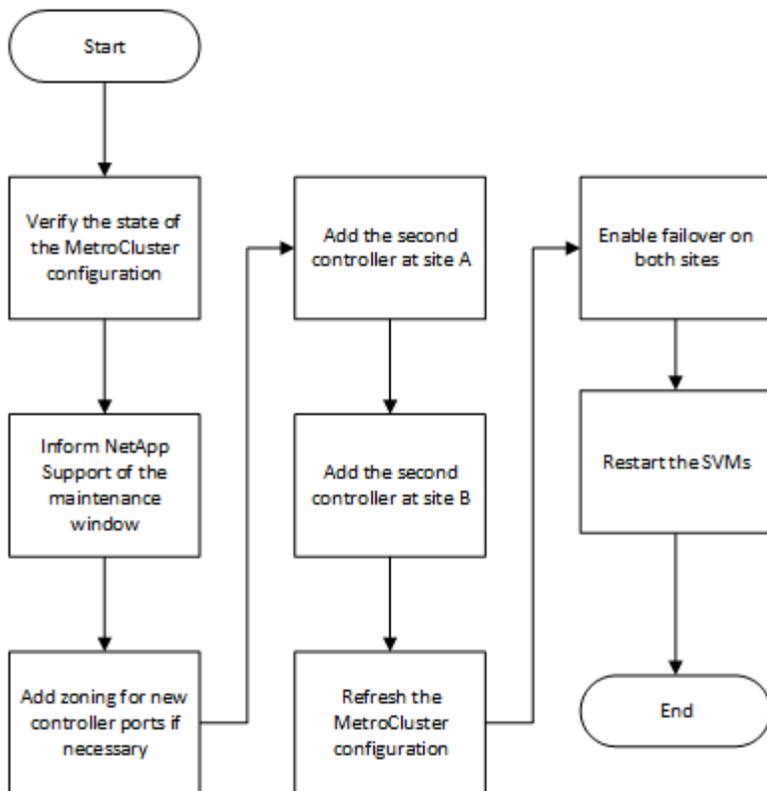
- Both sites must be expanded equally.

A MetroCluster configuration cannot consist of an uneven number of nodes.

- This procedure can take over an hour per site, with additional time for tasks such as initializing the disks and netbooting the new nodes.

The time to initialize the disks depends on the size of the disks.

- This procedure uses the following workflow:



## Enable console logging

Enable console logging on your devices before performing this task.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Verifying the state of the MetroCluster configuration

You should identify the existing controllers and confirm the disaster recovery (DR) relationships between them, that the controllers are in normal mode, and that the aggregates are mirrored.

### Steps

1. Display the details of the nodes in the MetroCluster configuration from any node in the configuration:

```
metrocluster node show -fields node,dr-partner,dr-partner-systemid
```

The following output shows that this MetroCluster configuration has a single DR group and one node in each cluster.

```
cluster_A::> metrocluster node show -fields node,dr-partner,dr-partner-
systemid

dr-group-id cluster node dr-partner dr-partner-
systemid

1 cluster_A controller_A_1 controller_B_1 536946192
1 cluster_B controller_B_1 controller_A_1 536946165
2 entries were displayed.
```

2. Display the state of the MetroCluster configuration:

```
metrocluster show
```

The following output shows that the existing nodes in the MetroCluster configuration are in normal mode:

```
cluster_A::> metrocluster show

Configuration: two-node-fabric

Cluster Entry Name State

Local: cluster_A Configuration State configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-
disaster
Remote: controller_B_1_siteB Configuration State configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-
disaster
```

3. Check the state of the aggregates on each node in the MetroCluster configuration:

```
storage aggregate show
```

The following output shows that the aggregates on cluster\_A are online and mirrored:

```
cluster_A::> storage aggregate show
```

| Aggregate<br>RAID Status | Size             | Available | Used% | State  | #Vols | Nodes |
|--------------------------|------------------|-----------|-------|--------|-------|-------|
| -----                    | -----            | -----     | ----- | -----  | ----- | ----- |
| aggr0_controller_A_1_0   | 1.38TB           | 68.63GB   | 95%   | online | 1     |       |
| controller_A_1           | raid_dp,mirrored |           |       |        |       |       |
| controller_A_1_aggr1     | 4.15TB           | 4.14TB    | 0%    | online | 2     |       |
| controller_A_1           | raid_dp,mirrored |           |       |        |       |       |
| controller_A_1_aggr2     | 4.15TB           | 4.14TB    | 0%    | online | 1     |       |
| controller_A_1           | raid_dp,mirrored |           |       |        |       |       |

3 entries were displayed.

```
cluster_A::>
```

## Sending a custom AutoSupport message before adding nodes to the MetroCluster configuration

You should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

### About this task

This task must be performed on each MetroCluster site.

### Steps

1. Log in to the cluster at Site\_A.
2. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-
window-in-hours
```

The `maintenance-window-in-hours` parameter specifies the length of the maintenance window and can be a maximum of 72 hours. If you complete the maintenance before the time has elapsed, you can issue the following command to indicate that the maintenance period has ended:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

3. Repeat this step on the partner site.

## Zoning for the new controller ports when adding a controller module in a fabric-attached MetroCluster configuration

The FC switch zoning must accommodate the new controller connections. If you used the NetApp-supplied reference configuration files (RCFs) to configure your switches, the

zoning is preconfigured and you do not need to make any changes.

If you manually configured your FC switches, you must ensure that the zoning is correct for the initiator connections from the new controller modules. See the sections on zoning in [Fabric-attached MetroCluster installation and configuration](#).

## Add a new controller module to each cluster

### Adding a new controller module to each cluster

You must add a new controller module to each site, creating an HA pair in each site. This is a multistep process involving both hardware and software changes that must be performed in the proper order at each site.

#### About this task

- The new controller module must be received from NetApp as part of the upgrade kit.

You should verify that PCIe cards in the new controller module are compatible and supported by the new controller module.

#### [NetApp Hardware Universe](#)

- Your system must have an empty slot available for the new controller module when upgrading to a single-chassis HA pair (an HA pair in which both controller modules reside in the same chassis).



This configuration is not supported on all systems. Platforms with single chassis configurations that are supported in ONTAP 9 are AFF A300, FAS8200, FAS8300, AFF A400, AFF80xx, FAS8020, FAS8060, FAS8080, and FAS9000.

- You must have rack space and cables for the new controller module when upgrading to a dual-chassis HA pair (an HA pair in which the controller modules reside in separate chassis).

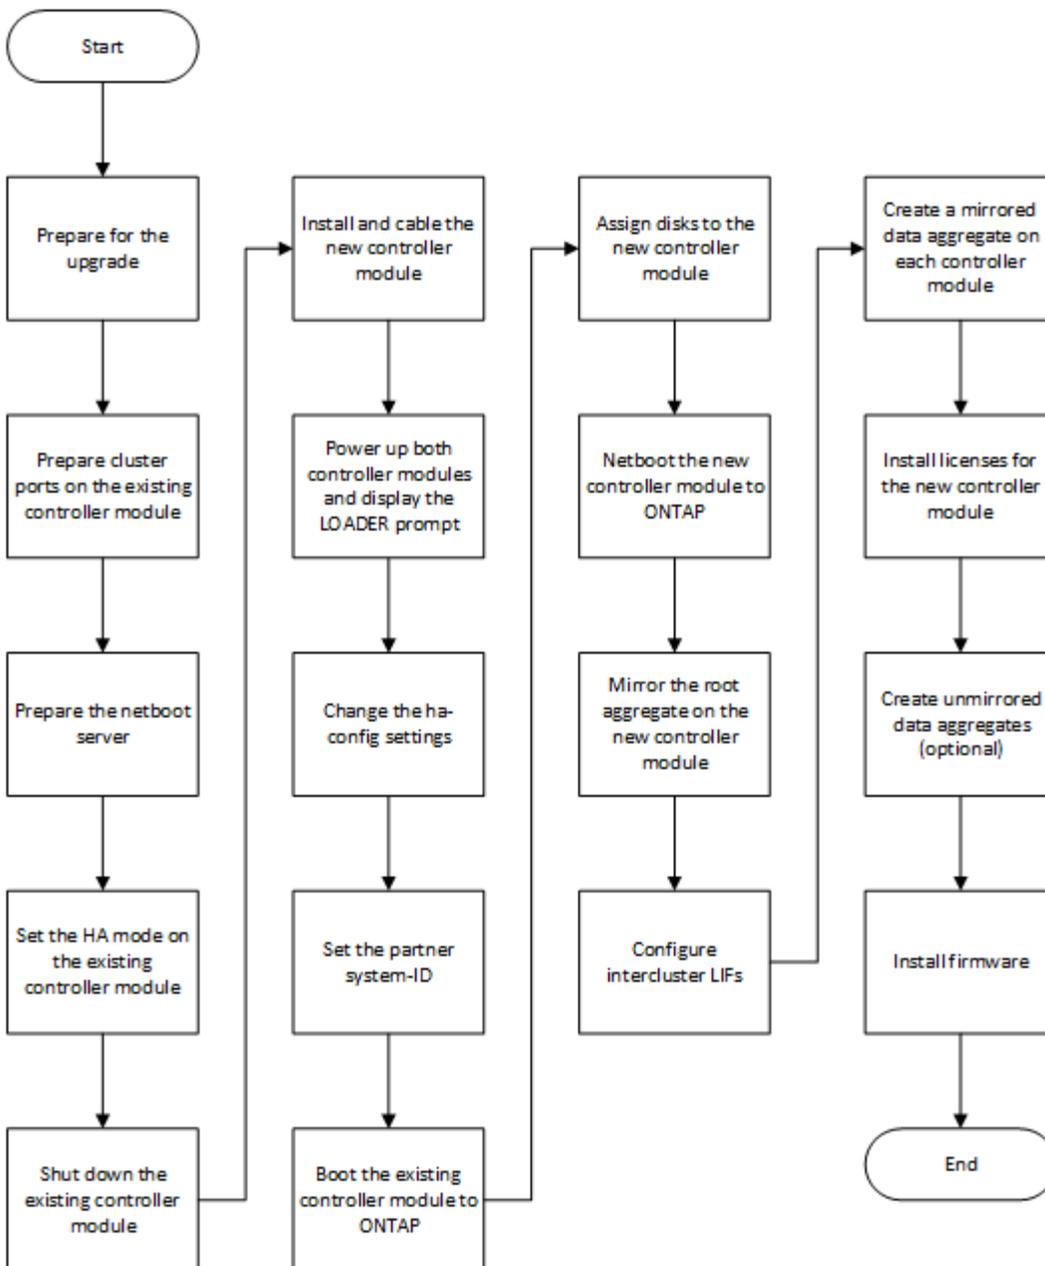


This configuration is not supported on all systems.

- You must connect each controller module to the management network through its e0a port or, if your system has one, you can connect to the e0M port as the management port.
- These tasks must be repeated at each site.
- The preexisting controller modules are referred to as the *existing* controller modules.

The examples in this procedure have the console prompt `existing_ctlr>`.

- The controller modules that are being added are referred to as the *new* controller modules; the examples in this procedure have the console prompt `new_ctlr>`.
- This task uses the following workflow:



## Preparing for the upgrade

Before upgrading to an HA pair, you must verify that your system meets all requirements and that you have all of the necessary information.

### Steps

1. Identify unassigned disks or spare disks that you can assign to the new controller module using the following commands:
  - ° `storage disk show -container-type spare`
  - ° `storage disk show -container-type unassigned`
2. Complete the following substeps:
  - a. Determine where the aggregates for the existing node are located:

```
storage aggregate show
```

- b. If disk ownership automatic assignment is on, turn it off:

```
storage disk option modify -node node_name -autoassign off
```

- c. Remove ownership on disks that do not have aggregates on them:

```
storage disk removeowner disk_name
```

- d. Repeat the previous step for as many disks as you need for the new node.

3. Verify that you have cables ready for the following connections:

- Cluster connections

If you are creating a two-node switchless cluster, you require two cables to connect the controller modules. Otherwise, you require a minimum of four cables, two for each controller module connection to the cluster-network switch. Other systems (like the 80xx series) have defaults of either four or six cluster connections.

- HA interconnect connections, if the system is in a dual-chassis HA pair

4. Verify that you have a serial port console available for the controller modules.

5. Verify that your environment meets the site and system requirements.

[NetApp Hardware Universe](#)

6. Gather all of the IP addresses and other network parameters for the new controller module.

### Clearing the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing configuration.

#### Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

## Preparing cluster ports on an existing controller module

Before installing a new controller module, you must configure cluster ports on the existing controller module so that the cluster ports can provide cluster communication with the new controller module.

### About this task

If you are creating a two-node switchless cluster (with no cluster network switches), you must enable the switchless cluster networking mode.

For detailed information about port, LIF, and network configuration in ONTAP, see [Network Management](#).

### Steps

1. Determine which ports should be used as the node's cluster ports.

For a list of the default port roles for your platform, see the [Hardware Universe](#)

The *Installation and Setup Instructions* for your platform on the NetApp Support Site contains information about the ports for cluster network connections.

2. For each cluster port, identify the port roles:

```
network port show
```

In the following example, ports "e0a", "e0b", "e0c", and "e0d" must be changed to cluster ports:

```
cluster_A::> network port show
```

```
Node: controller_A_1
```

```
Speed(Mbps) Health
```

| Port | IPspace | Broadcast Domain | Link | MTU  | Admin/Oper | Status  |
|------|---------|------------------|------|------|------------|---------|
| e0M  | Default | mgmt_bd_1500     | up   | 1500 | auto/1000  | healthy |
| e0a  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e0b  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e0c  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e0d  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e0i  | Default | Default          | down | 1500 | auto/10    | -       |
| e0j  | Default | Default          | down | 1500 | auto/10    | -       |
| e0k  | Default | Default          | down | 1500 | auto/10    | -       |
| e0l  | Default | Default          | down | 1500 | auto/10    | -       |
| e2a  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e2b  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e4a  | Default | Default          | up   | 1500 | auto/10000 | healthy |
| e4b  | Default | Default          | up   | 1500 | auto/10000 | healthy |

13 entries were displayed.

3. For any data LIF that is using a cluster port as the home-port or current-port, modify the LIF to use a data port as its home-port:

```
network interface modify
```

The following example changes the home port of a data LIF to a data port:

```
cluster1::> network interface modify -lif datalif1 -vserver vs1 -home
-port e1b
```

4. For each LIF that you modified, revert the LIF to its new home port:

```
network interface revert
```

The following example reverts the LIF “datalif1” to its new home port “e1b”:

```
cluster1::> network interface revert -lif datalif1 -vserver vs1
```

5. Remove any VLAN ports using cluster ports as member ports and ifgrps using cluster ports as member ports.

- a. Delete VLAN ports:

```
network port vlan delete -node node-name -vlan-name portid-vlandid
```

For example:

```
network port vlan delete -node node1 -vlan-name elc-80
```

b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node node-name -ifgrp interface-group-name
-port portid
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp ala -port e0d
```

c. Remove VLAN and interface group ports from broadcast domain::

```
network port broadcast-domain remove-ports -ipspace ipspace -broadcast
-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..
```

d. Modify interface group ports to use other physical ports as member as needed.:

```
ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id
```

6. Verify that the port roles have changed:

```
network port show
```

The following example shows that ports “e0a”, “e0b”, “e0c”, and “e0d” are now cluster ports:

```

Node: controller_A_1
Speed(Mbps) Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status

e0M Default mgmt_bd_1500 up 1500 auto/1000 healthy
e0a Cluster Cluster up 9000 auto/10000 healthy
e0b Cluster Cluster up 9000 auto/10000 healthy
e0c Cluster Cluster up 9000 auto/10000 healthy
e0d Cluster Cluster up 9000 auto/10000 healthy
e0i Default Default down 1500 auto/10 -
e0j Default Default down 1500 auto/10 -
e0k Default Default down 1500 auto/10 -
e0l Default Default down 1500 auto/10 -
e2a Default Default up 1500 auto/10000 healthy
e2b Default Default up 1500 auto/10000 healthy
e4a Default Default up 1500 auto/10000 healthy
e4b Default Default up 1500 auto/10000 healthy
13 entries were displayed.

```

7. Add the ports to the cluster broadcast domain:

```

broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster -ports
port-id, port-id, port-id...

```

For example:

```

broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster
-ports cluster1-01:e0a

```

8. If your system is part of a switched cluster, create cluster LIFs on the cluster ports: `network interface create`

The following example creates a cluster LIF on one of the node's cluster ports. The `-auto` parameter configures the LIF to use a link-local IP address.

```

cluster1::> network interface create -vserver Cluster -lif clus1 -role
cluster -home-node node0 -home-port e1a -auto true

```

9. If you are creating a two-node switchless cluster, enable the switchless cluster networking mode:

- a. Change to the advanced privilege level from either node:

```

set -privilege advanced

```

You can respond `y` when prompted whether you want to continue into advanced mode. The advanced

mode prompt appears (\*>).

- b. Enable the switchless cluster networking mode:

```
network options switchless-cluster modify -enabled true
```

- c. Return to the admin privilege level:

```
set -privilege admin
```



Cluster interface creation for the existing node in a two-node switchless cluster system is completed after cluster setup is completed through a netboot on the new controller module.

## Preparing the netboot server to download the image

When you are ready to prepare the netboot server, you must download the correct ONTAP netboot image from the NetApp Support Site to the netboot server and note the IP address.

### About this task

- You must be able to access an HTTP server from the system before and after adding the new controller module.
- You must have access to the NetApp Support Site to download the necessary system files for your platform and your version of ONTAP.

[NetApp Support Site](#)

- Both controller modules in the HA pair must run the same version of ONTAP.

### Steps

1. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `<ontap_version>_image.tgz` file on a web-accessible directory.

The `<ontap_version>_image.tgz` file is used for performing a netboot of your system.

2. Change to the web-accessible directory and verify that the files you need are available.

| For... | Then... |
|--------|---------|
|--------|---------|

|                                                                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems</p> | <p>Extract the contents of the <code>&lt;ontap_version&gt;_image.tgz</code> file to the target directory:</p> <pre>tar -zxvf &lt;ontap_version&gt;_image.tgz</pre> <p> If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</p> <p>Your directory listing should contain a netboot folder with a kernel file:</p> <pre>netboot/kernel</pre> |
| <p>All other systems</p>                                              | <p>Your directory listing should contain the following file:</p> <pre>&lt;ontap_version&gt;_image.tgz</pre> <p> There is no need to extract the file contents.</p>                                                                                                                                                                                                                      |

- Determine the IP address of the existing controller module.

This address is referred to later in this procedure as *ip-address-of-existing controller*.

- Ping *ip-address-of-existing controller* to verify that the IP address is reachable.

### Setting the HA mode on the existing controller module

You must use the storage failover modify command to set the mode on the existing controller module. The mode value is enabled later, after you reboot the controller module.

#### Steps

- Set the mode to HA:

```
storage failover modify -mode ha -node existing_node_name
```

### Shutting down the existing controller module

You must perform a clean shutdown of the existing controller module to verify that all of the data has been written to disk. You must also disconnect the power supplies.

#### About this task



You must perform a clean system shutdown before replacing the system components to avoid losing unwritten data in the NVRAM or NVMEM.

## Steps

1. Halt the node from the existing controller module prompt:

```
halt local -inhibit-takeover true
```

If you are prompted to continue the halt procedure, enter `y` when prompted, and then wait until the system stops at the LOADER prompt.

In an 80xx system, the NVRAM LED is located on the controller module to the right of the network ports, marked with a battery symbol.

This LED blinks if there is unwritten data in the NVRAM. If this LED is flashing amber after you enter the halt command, you need to reboot your system and try halting it again.

2. If you are not already grounded, properly ground yourself.
3. Turn off the power supplies and disconnect the power, using the correct method for your system and power-supply type:

| If your system uses... | Then...                                                                        |
|------------------------|--------------------------------------------------------------------------------|
| AC power supplies      | Unplug the power cords from the power source, and then remove the power cords. |
| DC power supplies      | Remove the power at the DC source, and then remove the DC wires, if necessary. |

## Install and cable the new controller module

### Installing and cabling the new controller module

You must physically install the new controller module in the chassis, and then cable it.

## Steps

1. If you have an I/O expansion module (IOXM) in your system and are creating a single-chassis HA pair, you must uncable and remove the IOXM.

You can then use the empty bay for the new controller module. However, the new configuration will not have the extra I/O provided by the IOXM.

2. Physically install the new controller module and, if necessary, install additional fans:

| If you are adding a controller module... | Then perform these steps... |
|------------------------------------------|-----------------------------|
|------------------------------------------|-----------------------------|

|                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>To an empty bay to create a single-chassis HA pair and the system belongs to one of the following platforms:</p>                                                                                                                        | <ol style="list-style-type: none"> <li>a. Remove the blank plate in the rear of the chassis that covers the empty bay that will contain the new controller module.</li> <li>b. Gently push the controller module halfway into the chassis.</li> </ol> <p>To prevent the controller module from automatically booting, do not fully seat it in the chassis until later in this procedure.</p> |
| <p>In a separate chassis from its HA partner to create a dual-chassis HA pair when the existing configuration is in a controller-IOX module configuration.</p> <ul style="list-style-type: none"> <li>• FAS8200</li> <li>• 80xx</li> </ul> | <p>Install the new system in the rack or system cabinet.</p>                                                                                                                                                                                                                                                                                                                                 |

3. Cable the cluster network connections, as necessary:

- a. Identify the ports on the controller module for the cluster connections.

[AFF A320 systems: Installation and setup](#)

[AFF A220/FAS2700 Systems Installation and Setup Instructions](#)

[AFF A800 Systems Installation and Setup Instructions](#)

[AFF A300 Systems Installation and Setup Instructions](#)

[FAS8200 Systems Installation and Setup Instructions](#)

- b. If you are configuring a switched cluster, identify the ports that you will use on the cluster network switches.

See the [Clustered Data ONTAP Switch Setup Guide for Cisco Switches](#), [NetApp 10G Cluster-Mode Switch Installation Guide](#) or [NetApp 1G Cluster-Mode Switch Installation Guide](#), depending on what switches you are using.

- c. Connect cables to the cluster ports:

| If the cluster is...          | Then...                                                                                                                               |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| A two-node switchless cluster | Directly connect the cluster ports on the existing controller module to the corresponding cluster ports on the new controller module. |
| A switched cluster            | Connect the cluster ports on each controller to the ports on the cluster network switches identified in Substep b.                    |

### Cabling the new controller module's FC-VI and HBA ports to the FC switches

The new controller module's FC-VI ports and HBAs (host bus adapters) must be cabled to the site FC switches.

#### Steps

1. Cable the FC-VI ports and HBA ports, using the table for your configuration and switch model.
  - [Port assignments for FC switches](#)
  - [Port assignments for systems using two initiator ports](#)

### Cabling the new controller module's cluster peering connections

You must cable the new controller module to the cluster peering network so that it has connectivity with the cluster on the partner site.

#### About this task

At least two ports on each controller module should be used for cluster peering.

The recommended minimum bandwidth for the ports and network connectivity is 1 GbE.

#### Steps

1. Identify and cable at least two ports for cluster peering and verify they have network connectivity with the partner cluster.

### Powering up both controller modules and displaying the LOADER prompt

You power up the existing controller module and the new controller module to display the LOADER prompt.

#### Steps

Power up the controller modules and interrupt the boot process, using the steps for your configuration:

| If the controller modules are... | Then... |
|----------------------------------|---------|
|                                  |         |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| In the same chassis | <p>a. Verify that the new controller module is <b>not</b> fully inserted into the bay.</p> <p>The existing controller module should be fully inserted into the bay because it was never removed from the chassis, but the new controller module should not be.</p> <p>b. Connect the power and turn on the power supplies so that the existing controller module receives power.</p> <p>c. Interrupt the boot process on the existing controller module by pressing Ctrl-C.</p> <p>d. Push the new controller module firmly into the bay.</p> <p>When fully seated, the new controller module receives power and automatically boots.</p> <p>e. Interrupt the boot process by pressing Ctrl-C.</p> <p>f. Tighten the thumbscrew on the cam handle, if present.</p> <p>g. Install the cable management device, if present.</p> <p>h. Bind the cables to the cable management device with the hook and loop strap.</p> |
| In separate chassis | <p>a. Turn on the power supplies on the existing controller module.</p> <p>b. Interrupt the boot process by pressing Ctrl-C.</p> <p>c. Repeat these steps for the new controller module</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

Each controller module should display the LOADER prompt (LOADER>, LOADER-A>, or LOADER-B>).



If there is no LOADER prompt, record the error message. If the system displays the boot menu, reboot and attempt to interrupt the boot process again.

### Changing the ha-config setting on the existing and new controller modules

When you expand a MetroCluster configuration, you must update the ha-config setting of the existing controller module and the new controller module. You must also determine the system ID of the new controller module.

#### About this task

This task is performed in Maintenance mode on both the existing and new controller modules.

#### Steps

1. Change the ha-config setting of the existing controller module:
  - a. Display the ha-config setting of the existing controller module and chassis:

```
ha-config show
```

The ha-config setting is “mcc-2n” for all components because the controller module was in a two-node MetroCluster configuration.

- b. Change the ha-config setting of the existing controller module to “mcc”:

```
ha-config modify controller mcc
```

- c. Change the ha-config setting of the existing chassis to “mcc”:

```
ha-config modify chassis mcc
```

- d. Retrieve the system ID for the existing controller module:

```
sysconfig
```

Note the system ID. You need it when you set the partner ID on the new controller module.

- e. Exit Maintenance mode to return to the LOADER prompt:

```
halt
```

2. Change the ha-config setting and retrieve the system ID of the new controller module:

- a. If the new controller module is not already in Maintenance mode, boot it to Maintenance mode:

```
boot_ontap maint
```

- b. Change the ha-config setting of the new controller module to “mcc”:

```
ha-config modify controller mcc
```

- c. Change the ha-config setting of the new chassis to mcc:

```
ha-config modify chassis mcc
```

- d. Retrieve the system ID for the new controller module:

```
sysconfig
```

Note the system ID. You need it when you set the partner ID and assign disks to the new controller module.

- e. Exit Maintenance mode to return to the LOADER prompt:

```
halt
```

### Setting the partner system ID for both controller modules

You must set the partner system ID on both controller modules so that they can form an HA pair.

#### About this task

This task is performed with both controller modules at the LOADER prompt.

#### Steps

1. On the existing controller module, set the partner system ID to that of the new controller module:

```
setenv partner-sysid sysID_of_new_controller
```

2. On the new controller module, set the partner system ID to that of the existing controller module:

```
setenv partner-sysid sysID_of_existing_controller
```

## Booting the existing controller module

You must boot the existing controller module to ONTAP.

### Steps

1. At the LOADER prompt, boot the existing controller module to ONTAP:

```
boot_ontap
```

## Assigning disks to the new controller module

Before you complete the configuration of the new controller module through netboot, you must assign disks to it.

### About this task

You must have made sure that there are enough spares, unassigned disks, or assigned disks that are not part of an existing aggregate.

### [Preparing for the upgrade](#)

These steps are performed on the existing controller module.

### Steps

1. Assign the root disk to the new controller module:

```
storage disk assign -disk disk_name -sysid new_controller_sysID -force true
```

If your platform model uses the Advanced Drive Partitioning (ADP) feature, you must include the `-root true` parameter:

```
storage disk assign -disk disk_name -root true -sysid new_controller_sysID -force true
```

2. Assign the remaining required disks to the new controller module by entering the following command for each disk:

```
storage disk assign -disk disk_name -sysid new_controller_sysID -force true
```

3. Verify that the disk assignments are correct:

```
storage disk show -partitionownership*
```



Ensure that you have assigned all disks that you intend to assign to the new node.

## Netbooting and setting up ONTAP on the new controller module

You must perform a specific sequence of steps to netboot and install the ONTAP operating system on the new controller module when adding controller modules to an existing MetroCluster configuration.

### About this task

- This task starts at the LOADER prompt of the new controller module.
- This task includes initializing disks.

The amount of time you need to initialize the disks depends on the size of the disks.

- The system automatically assigns two disks to the new controller module.

### [Disk and aggregate management](#)

### Steps

1. At the LOADER prompt, configure the IP address of the new controller module based on DHCP availability:

| If DHCP is... | Then enter the following command...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Available     | <b>ifconfig e0M -auto</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Not available | <pre>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></pre> <p><i>filer_addr</i> is the IP address of the storage system.</p> <p><i>netmask</i> is the network mask of the storage system.</p> <p><i>gateway</i> is the gateway for the storage system.</p> <p><i>dns_addr</i> is the IP address of a name server on your network.</p> <p><i>dns_domain</i> is the Domain Name System (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.</p> <div style="border-left: 1px solid #ccc; padding-left: 10px;"> Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.</div> |

2. At the LOADER prompt, netboot the new node:

| For... | Issue this command... |
|--------|-----------------------|
|--------|-----------------------|

|                                                                |                                                                                          |
|----------------------------------------------------------------|------------------------------------------------------------------------------------------|
| FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems | netboot http://web_server_ip/path_to_web-accessible_directory/netboot/kernel             |
| All other systems                                              | netboot \http://web_server_ip/path_to_web-accessible_directory/<ontap_version>_image.tgz |

The `path_to_the_web-accessible_directory` is the location of the downloaded `<ontap_version>_image.tgz` file.

3. Select the **Install new software first** option from the displayed menu.

This menu option downloads and installs the new ONTAP image to the boot device.

- You should enter “y” when prompted with the message that this procedure is not supported for nondisruptive upgrade on an HA pair.
- You should enter “y” when warned that this process replaces the existing ONTAP software with new software.
- You should enter the path as follows when prompted for the URL of the image.tgz file:

```
http://path_to_the_web-accessible_directory/image.tgz
```

4. Enter “y” when prompted regarding nondisruptive upgrade or replacement of the software.

5. Enter the path to the image.tgz file when prompted for the URL of the package.

```
What is the URL for the package? `http://path_to_web-accessible_directory/image.tgz`
```

6. Enter “n” to skip the backup recovery when prompted to restore the backup configuration.

```

* Restore Backup Configuration *
* This procedure only applies to storage controllers that *
* are configured as an HA pair. *
* *
* Choose Yes to restore the "varfs" backup configuration *
* from the SSH server. Refer to the Boot Device Replacement *
* guide for more details. *
* Choose No to skip the backup recovery and return to the *
* boot menu. *

Do you want to restore the backup configuration
now? {y|n} `n`
```

7. Enter “y” when prompted to reboot now.

```
The node must be rebooted to start using the newly installed software.
Do you want to
reboot now? {y|n} `y`
```

8. If necessary, select the option to **Clean configuration and initialize all disks** after the node has booted.

Because you are configuring a new controller module and the new controller module’s disks are empty, you can respond “y” when the system warns you that this will erase all disks.



The amount of time needed to initialize disks depends on the size of your disks and configuration.

9. After the disks are initialized and the Cluster Setup wizard starts, set up the node:

Enter the node management LIF information on the console.

10. Log in to the node, and enter the `cluster setup` and then enter “join” when prompted to join the cluster.

```
Do you want to create a new cluster or join an existing cluster?
{create, join}: `join`
```

11. Respond to the remaining prompts as appropriate for your site.

The [Setup ONTAP](#) for your version of ONTAP contains additional details.

12. If the system is in a two-node switchless cluster configuration, create the cluster interfaces on the existing node using the network interface create command to create cluster LIFs on the cluster ports.

The following is an example command for creating a cluster LIF on one of the node’s cluster ports. The `-auto` parameter configures the LIF to use a link-local IP address.

```
cluster_A::> network interface create -vserver Cluster -lif clus1 -role
cluster -home-node node_A_1 -home-port e1a -auto true
```

13. After setup is complete, verify that the node is healthy and eligible to participate in the cluster:

```
cluster show
```

The following example shows a cluster after the second node (cluster1-02) has been joined to it:

```
cluster_A::> cluster show
Node Health Eligibility

node_A_1 true true
node_A_2 true true
```

You can access the Cluster Setup wizard to change any of the values you entered for the admin storage virtual machine (SVM) or node SVM by using the cluster setup command.

14. Confirm that you have four ports configured as cluster interconnects:

```
network port show
```

The following example shows output for two controller modules in cluster\_A:

```
cluster_A::> network port show

(Mbps)
Node Port IPspace Broadcast Domain Link MTU Admin/Oper Speed

node_A_1
 **e0a Cluster Cluster up 9000
auto/1000
 e0b Cluster Cluster up 9000
auto/1000**
 e0c Default Default up 1500 auto/1000
 e0d Default Default up 1500 auto/1000
 e0e Default Default up 1500 auto/1000
 e0f Default Default up 1500 auto/1000
 e0g Default Default up 1500 auto/1000
node_A_2
 **e0a Cluster Cluster up 9000
auto/1000
 e0b Cluster Cluster up 9000
auto/1000**
 e0c Default Default up 1500 auto/1000
 e0d Default Default up 1500 auto/1000
 e0e Default Default up 1500 auto/1000
 e0f Default Default up 1500 auto/1000
 e0g Default Default up 1500 auto/1000
14 entries were displayed.
```

## Mirroring the root aggregate on the new controller

You must mirror the root aggregate to provide data protection when you are adding a controller to a MetroCluster configuration.

This task must be performed on the new controller module.

1. Mirror the root aggregate:

```
storage aggregate mirror aggr_name
```

The following command mirrors the root aggregate for controller\_A\_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote MetroCluster site.

## Configure intercluster LIFs

Learn how to configure intercluster LIFs on dedicated and shared ports.

### **Configure intercluster LIFs on dedicated ports**

You can configure intercluster LIFs on dedicated ports to increase the available bandwidth for replication traffic.

#### **Steps**

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```

cluster01::> network port show

```

|              |      |         |                  |      |      | Speed |
|--------------|------|---------|------------------|------|------|-------|
| (Mbps)       |      |         |                  |      |      |       |
| Node         | Port | IPspace | Broadcast Domain | Link | MTU  |       |
| Admin/Oper   |      |         |                  |      |      |       |
| -----        |      |         |                  |      |      | ----- |
| cluster01-01 |      |         |                  |      |      |       |
|              | e0a  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0b  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0c  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0d  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0e  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0f  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
| cluster01-02 |      |         |                  |      |      |       |
|              | e0a  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0b  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0c  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0d  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0e  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0f  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |

2. Determine which ports are available to dedicate to intercluster communication:

```
network interface show -fields home-port,curr-port
```

For complete command syntax, see the man page.

The following example shows that ports "e0e" and "e0f" have not been assigned LIFs:

```

cluster01::> network interface show -fields home-port,curr-port
vserver lif home-port curr-port

Cluster cluster01-01_clus1 e0a e0a
Cluster cluster01-01_clus2 e0b e0b
Cluster cluster01-02_clus1 e0a e0a
Cluster cluster01-02_clus2 e0b e0b
cluster01
 cluster_mgmt e0c e0c
cluster01
 cluster01-01_mgmt1 e0c e0c
cluster01
 cluster01-02_mgmt1 e0c e0c

```

### 3. Create a failover group for the dedicated ports:

```

network interface failover-groups create -vserver <system_SVM> -failover
-group <failover_group> -targets <physical_or_logical_ports>

```

The following example assigns ports "e0e" and "e0f" to the failover group "intercluster01" on the system SVM "cluster01":

```

cluster01::> network interface failover-groups create -vserver
cluster01 -failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f

```

### 4. Verify that the failover group was created:

```

network interface failover-groups show

```

For complete command syntax, see the man page.

```

cluster01::> network interface failover-groups show
 Failover
Vserver Group Targets

Cluster
 Cluster
 cluster01-01:e0a, cluster01-
01:e0b,
 cluster01-02:e0a, cluster01-02:e0b
cluster01
 Default
 cluster01-01:e0c, cluster01-
01:e0d,
 cluster01-02:e0c, cluster01-
02:e0d,
 cluster01-01:e0e, cluster01-01:e0f
 cluster01-02:e0e, cluster01-02:e0f
 intercluster01
 cluster01-01:e0e, cluster01-01:e0f
 cluster01-02:e0e, cluster01-02:e0f

```

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

| ONTAP version   | Command                                                                                                                                                                                                        |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 9.6 and later   | network interface create -vserver <system_SVM> -lif <LIF_name> -service-policy default-intercluster -home-node <node> -home-port <port> -address <port_IP> -netmask <netmask> -failover-group <failover_group> |
| 9.5 and earlier | network interface create -vserver system_SVM -lif <LIF_name> -role intercluster -home-node <node> -home-port <port> -address <port_IP> -netmask <netmask> -failover-group <failover_group>                     |

For complete command syntax, see the man page.

The following example creates intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" in the failover group "intercluster01":

```

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0 -failover-group intercluster01

```

6. Verify that the intercluster LIFs were created:

**In ONTAP 9.6 and later:**

```
network interface show -service-policy default-intercluster
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster
```

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-
intercluster

 Logical Status Network Current
Current Is
Vserver Interface Admin/Oper Address/Mask Node
Port Home

cluster01
 cluster01_icl01
 up/up 192.168.1.201/24 cluster01-01
e0e true
 cluster01_icl02
 up/up 192.168.1.202/24 cluster01-02
e0f true

```

7. Verify that the intercluster LIFs are redundant:

**In ONTAP 9.6 and later:**

```
network interface show -service-policy default-intercluster -failover
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" on the SVM "e0e" port will fail over to the "e0f" port.

```
cluster01::> network interface show -service-policy default-
intercluster -failover
 Logical Home Failover
Failover
Vserver Interface Node:Port Policy Group
----- -----
cluster01
 cluster01_icl01 cluster01-01:e0e local-only
intercluster01
 Failover Targets: cluster01-01:e0e,
 cluster01-01:e0f
 cluster01_icl02 cluster01-02:e0e local-only
intercluster01
 Failover Targets: cluster01-02:e0e,
 cluster01-02:e0f
```

### Configure intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network to reduce the number of ports you need for intercluster networking.

#### Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```
cluster01::> network port show
```

|              |       |         |                  |       |       | Speed |
|--------------|-------|---------|------------------|-------|-------|-------|
| (Mbps)       |       |         |                  |       |       |       |
| Node         | Port  | IPspace | Broadcast Domain | Link  | MTU   |       |
| Admin/Oper   |       |         |                  |       |       |       |
| -----        | ----- | -----   | -----            | ----- | ----- | ----- |
| cluster01-01 |       |         |                  |       |       |       |
|              | e0a   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0b   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0c   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0d   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
| cluster01-02 |       |         |                  |       |       |       |
|              | e0a   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0b   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0c   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0d   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |

## 2. Create intercluster LIFs on the system SVM:

### In ONTAP 9.6 and later:

```
network interface create -vserver <system_SVM> -lif <LIF_name> -service
-policy default-intercluster -home-node <node> -home-port <port> -address
<port_IP> -netmask <netmask>
```

### In ONTAP 9.5 and earlier:

```
network interface create -vserver <system_SVM> -lif <LIF_name> -role
intercluster -home-node <node> -home-port <port> -address <port_IP>
-netmask <netmask>
```

For complete command syntax, see the man page.

The following example creates intercluster LIFs `cluster01_icl01` and `cluster01_icl02`:

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0c
-address 192.168.1.201
-netmask 255.255.255.0
```

```
cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0c
-address 192.168.1.202
-netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

**In ONTAP 9.6 and later:**

```
network interface show -service-policy default-intercluster
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster
```

For complete command syntax, see the man page.

```
cluster01::> network interface show -service-policy default-
intercluster
```

| Current Is | Logical         | Status     | Network          | Current      |
|------------|-----------------|------------|------------------|--------------|
| Vserver    | Interface       | Admin/Oper | Address/Mask     | Node         |
| Port       | Home            |            |                  |              |
| -----      | -----           | -----      | -----            | -----        |
| -----      | -----           | -----      | -----            | -----        |
| cluster01  | cluster01_icl01 | up/up      | 192.168.1.201/24 | cluster01-01 |
| e0c        | true            |            |                  |              |
|            | cluster01_icl02 | up/up      | 192.168.1.202/24 | cluster01-02 |
| e0c        | true            |            |                  |              |

4. Verify that the intercluster LIFs are redundant:

**In ONTAP 9.6 and later:**

```
network interface show -service-policy default-intercluster -failover
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" on the "e0c" port will fail over to the "e0d" port.

```
cluster01::> network interface show -service-policy default-
intercluster -failover
 Logical Home Failover
Failover
Vserver Interface Node:Port Policy Group
----- -
cluster01
 cluster01_icl01 cluster01-01:e0c local-only
192.168.1.201/24
 Failover Targets: cluster01-01:e0c,
 cluster01-01:e0d
 cluster01_icl02 cluster01-02:e0c local-only
192.168.1.201/24
 Failover Targets: cluster01-02:e0c,
 cluster01-02:e0d
```

## Create a mirrored data aggregate on each MetroCluster FC node

You must create a mirrored data aggregate on each node in the DR group.

### About this task

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can make sure that the correct drive type is selected.
- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.

In systems using ADP, aggregates are created using partitions in which each drive is partitioned in to P1, P2 and P3 partitions.

- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.

### [Disk and aggregate management](#)

- Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.



It's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space can be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices might have a negative impact on performance.

## Steps

1. Display a list of available spares:

```
storage disk show -spare -owner <node_name>
```

2. Create the aggregate:

```
storage aggregate create -mirror true
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To verify that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include



In the minimum supported configuration, in which a limited number of drives are available, you must use the `force-small-aggregate` option to allow the creation of a three disk RAID-DP aggregate.

- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create` man page.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10
-node node_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_node_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate <aggregate-name>
```

### Installing licenses for the new controller module

You must add licenses for the new controller module for any ONTAP services that require standard (node-locked) licenses. For features with standard licenses, each node in the cluster must have its own key for the feature.

For detailed information about licensing, see the knowledgebase article 3013749: Data ONTAP 8.2 Licensing Overview and References on the NetApp Support Site and the *System Administration Reference*.

#### Steps

1. If necessary, obtain license keys for the new node on the NetApp Support Site in the My Support section under Software licenses.

For further information on license replacements, see the Knowledge Base article [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#).

2. Issue the following command to install each license key:

```
system license add -license-code license_key
```

The *license\_key* is 28 digits in length.

3. Repeat this step for each required standard (node-locked) license.

### Creating unmirrored data aggregates

You can optionally create unmirrored data aggregates for data that does not require the redundant mirroring provided by MetroCluster configurations.

#### About this task

- Verify that you know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can verify that the correct drive type is selected.



In MetroCluster IP configurations, remote unmirrored aggregates are not accessible after a switchover



The unmirrored aggregates must be local to the node owning them.

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- *Disks and aggregates management* contains more information about mirroring aggregates.

## Steps

1. Install and cable the disk shelves that will contain the unmirrored aggregates.

You can use the procedures in the *Installation and Setup* documentation for your platform and disk shelves.

[ONTAP Hardware Systems Documentation](#)

2. Manually assign all disks on the new shelf to the appropriate node:

```
disk assign -disk <disk-id> -owner <owner-node-name>
```

3. Create the aggregate:

```
storage aggregate create
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To verify that the aggregate is created on a specific node, you should use the `-node` parameter or specify drives that are owned by that node.

Verify that you are only including drives on the unmirrored shelf to the aggregate.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include
- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create man` page.

The following command creates a unmirrored aggregate with 10 disks:

```
controller_A_1::> storage aggregate create aggr1_controller_A_1
-diskcount 10 -node controller_A_1
[Job 15] Job is queued: Create aggr1_controller_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```



You can also use the `-disklist` parameter in the command to specify the disks that you want to use for the aggregate.

4. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate <aggregate-name>
```

### Related information

[Disk and aggregate management](#)

### Installing the firmware after adding a controller module

After adding the controller module, you must install the latest firmware on the new controller module so that the controller module functions properly with ONTAP.

#### Steps

1. Download the most current version of firmware for your system and follow the instructions for downloading and installing the new firmware.

[NetApp Downloads: System Firmware and Diagnostics](#)

### Refreshing the MetroCluster configuration with new controllers

You must refresh the MetroCluster configuration when expanding it from a two-node configuration to a four-node configuration.

#### Steps

1. Refresh the MetroCluster configuration:

- a. Enter advanced privilege mode:

```
set -privilege advanced
```

- b. Refresh the MetroCluster configuration:

```
metrocluster configure -refresh true -allow-with-one-aggregate true
```

The following command refreshes the MetroCluster configuration on all of the nodes in the DR group that contains `controller_A_1`:

```
controller_A_1::*> metrocluster configure -refresh true -allow-with
-one-aggregate true
```

```
[Job 726] Job succeeded: Configure is successful.
```

c. Return to admin privilege mode:

```
set -privilege admin
```

2. Verify the networking status on site A:

```
network port show
```

The following example shows the network port usage on a four-node MetroCluster configuration:

```
cluster_A::> network port show
```

| Node           | Port | IPspace | Broadcast Domain | Link | MTU  | Speed (Mbps)<br>Admin/Oper |
|----------------|------|---------|------------------|------|------|----------------------------|
| controller_A_1 |      |         |                  |      |      |                            |
|                | e0a  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0b  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0c  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0d  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0e  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0f  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0g  | Default | Default          | up   | 1500 | auto/1000                  |
| controller_A_2 |      |         |                  |      |      |                            |
|                | e0a  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0b  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0c  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0d  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0e  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0f  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0g  | Default | Default          | up   | 1500 | auto/1000                  |

14 entries were displayed.

3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration.

a. Verify the configuration from site A:

```
metrocluster show
```

```

cluster_A::> metrocluster show
Cluster Entry Name State

Local: cluster_A Configuration state configured
Mode normal
AUSO Failure Domain auso-on-cluster-
disaster
Remote: cluster_B Configuration state configured
Mode normal
AUSO Failure Domain auso-on-cluster-
disaster

```

b. Verify the configuration from site B:

```
metrocluster show
```

```

cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
Mode normal
AUSO Failure Domain auso-on-cluster-
disaster
Remote: cluster_A Configuration state configured
Mode normal
AUSO Failure Domain auso-on-cluster-
disaster

```

c. Verify that the DR relationships have been created correctly:

```
metrocluster node show -fields dr-cluster,dr-auxiliary,node-object-
limit,automatic-uso,ha-partner,dr-partner
```

```

metrocluster node show -fields dr-cluster,dr-auxiliary,node-object-
limit,automatic-uso,ha-partner,dr-partner
dr-group-id cluster node ha-partner dr-cluster dr-partner
dr-auxiliary node-object-limit automatic-uso

2 cluster_A node_A_1 node_A_2 cluster_B node_B_1
node_B_2 on true
2 cluster_A node_A_2 node_A_1 cluster_B node_B_2
node_B_1 on true
2 cluster_B node_B_1 node_B_2 cluster_A node_A_1
node_A_2 on true
2 cluster_B node_B_2 node_B_1 cluster_A node_A_2
node_A_1 on true
4 entries were displayed.

```

## Enabling storage failover on both controller modules and enabling cluster HA

After adding new controller modules to the MetroCluster configuration, you must enable storage failover on both controller modules and separately enable cluster HA.

### Before you begin

The MetroCluster configuration must have previously been refreshed using the `metrocluster configure -refresh true` command.

### About this task

This task must be performed on each MetroCluster site.

### Steps

1. Enable storage failover:

```
storage failover modify -enabled true -node existing-node-name
```

The single command enables storage failover on both controller modules.

2. Verify that storage failover is enabled:

```
storage failover show
```

The output should be similar to the following:

| Node     | Partner  | Possible State | Description           |
|----------|----------|----------------|-----------------------|
| old-ctrl | new-ctrl | true           | Connected to new-ctrl |
| new-ctrl | old-ctrl | true           | Connected to old-ctrl |

2 entries were displayed.

### 3. Enable cluster HA:

```
cluster ha modify -configured true
```

Cluster high availability (HA) must be configured in a cluster if it contains only two nodes and it differs from the HA provided by storage failover.

## Restarting the SVMs

After expanding the MetroCluster configuration, you must restart the SVMs.

### Steps

#### 1. Identify the SVMs that need to be restarted:

```
metrocluster vserver show
```

This command shows the SVMs on both MetroCluster clusters.

#### 2. Restart the SVMs on the first cluster:

##### a. Enter advanced privilege mode, pressing **y** when prompted:

```
set -privilege advanced
```

##### b. Restart the SVMs:

```
vserver start -vserver SVM_name -force true
```

##### c. Return to admin privilege mode:

```
set -privilege admin
```

#### 3. Repeat the previous step on the partner cluster.

#### 4. Verify that the SVMs are in a healthy state:

```
metrocluster vserver show
```

## Expand a four-node MetroCluster FC configuration to an eight-node configuration

### Expanding a four-node MetroCluster FC configuration to an eight-node configuration

Expanding a four-node MetroCluster FC configuration to an eight-node MetroCluster FC configuration involves adding two controllers to each cluster to form a second HA pair at each MetroCluster site, and then running the MetroCluster FC configuration operation.

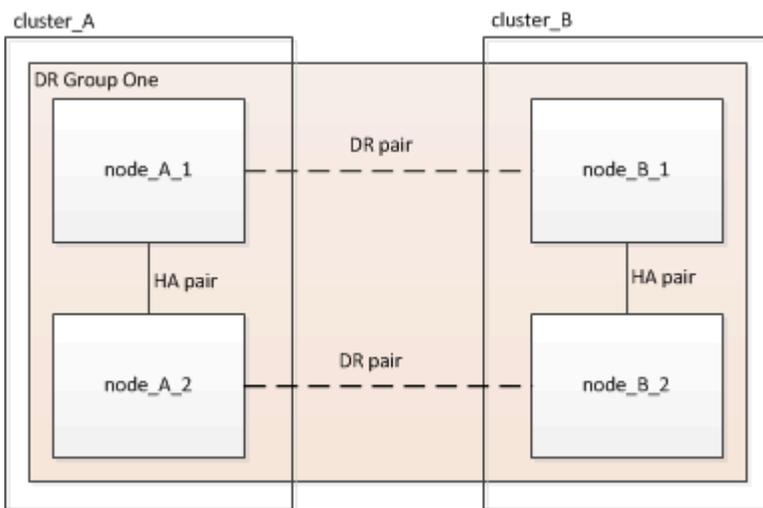
#### About this task

- The nodes must be running ONTAP 9 in a MetroCluster FC configuration.

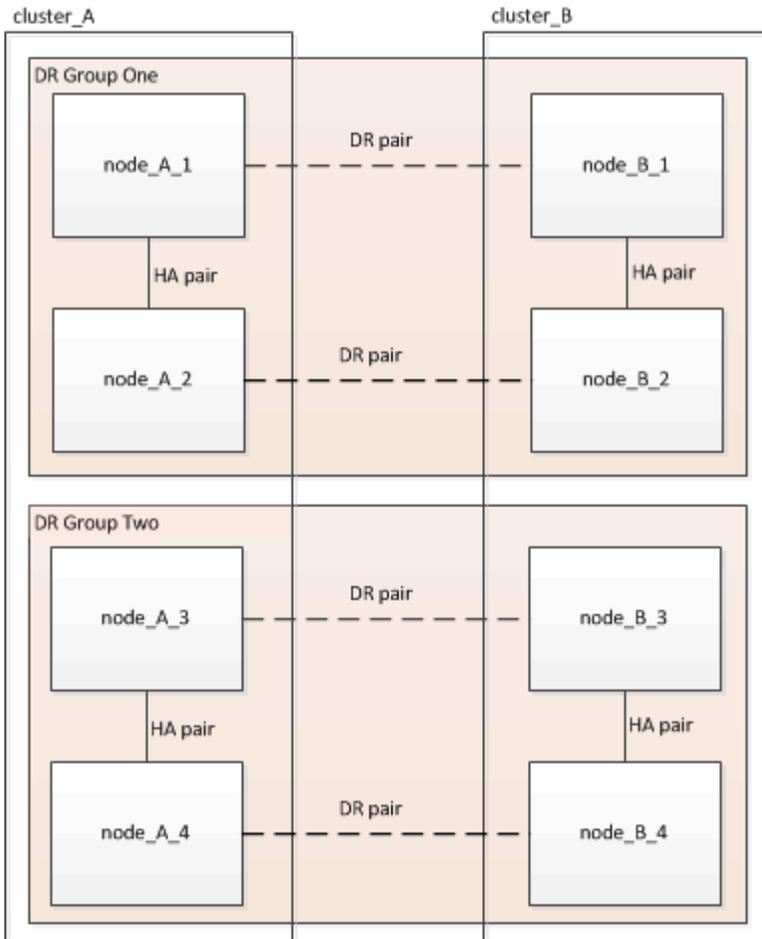
This procedure is not supported on earlier versions of ONTAP or in MetroCluster IP configurations.

- The existing MetroCluster FC configuration must be healthy.
- The equipment you are adding must be supported and meet all the requirements described in [Fabric-attached MetroCluster installation and configuration](#)
- You must have available FC switch ports to accommodate the new controllers and any new bridges.
- You need the admin password and access to an FTP or SCP server.
- This procedure applies only to MetroCluster FC configurations.
- This procedure is nondisruptive and takes approximately one day to complete (excluding rack and stack) when disks are zeroed.

Before performing this procedure, the MetroCluster FC configuration consists of four nodes, with one HA pair at each site:



At the conclusion of this procedure, the MetroCluster FC configuration consists of two HA pairs at each site:



Both sites must be expanded equally. A MetroCluster FC configuration cannot consist of an uneven number of nodes.

### Supported platform combinations when adding a second DR group

The following tables shows the supported platform combinations for eight-node MetroCluster FC configurations.



- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, if you have an eight-node configuration, all eight nodes must be running the same ONTAP version.
- The combinations in this table apply only to regular or permanent eight-node configurations.
- The platform combinations in this table **do not** apply if you are using the the transition or refresh procedures.
- All nodes in one DR group must be of the same type and configuration.

### Supported AFF and FAS MetroCluster FC expansion combinations

The following table shows the supported platform combinations for expanding an AFF or FAS system in a MetroCluster FC configuration:

| FAS and AFF           |          | Eight-node DR group 2 |          |         |          |         |          |         |          |
|-----------------------|----------|-----------------------|----------|---------|----------|---------|----------|---------|----------|
|                       |          | FAS8200               | AFF A300 | FAS8300 | AFF A400 | FAS9000 | AFF A700 | FAS9500 | AFF A900 |
| Eight-node DR group 1 | FAS8200  |                       |          |         |          |         |          |         |          |
|                       | AFF A300 |                       |          |         |          |         |          |         |          |
|                       | FAS8300  |                       |          |         |          |         |          |         |          |
|                       | AFF A400 |                       |          |         |          |         |          |         |          |
|                       | FAS9000  |                       |          |         |          |         |          |         |          |
|                       | AFF A700 |                       |          |         |          |         |          |         |          |
|                       | FAS9500  |                       |          |         |          |         |          |         |          |
|                       | AFF A900 |                       |          |         |          |         |          |         |          |

### Supported ASA MetroCluster FC expansion combinations

The following table shows the supported platform combinations for expanding an ASA system in a MetroCluster FC configuration:

| Eight-node DR group 1 | Eight-node DR group 2 | Supported? |
|-----------------------|-----------------------|------------|
| ASA A400              | ASA A400              | Yes        |
|                       | ASA A900              | No         |
| ASA A900              | ASA A400              | No         |
|                       | ASA A900              | Yes        |

### Enable console logging

Enable console logging on your devices before performing this task.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

### Determining the new cabling layout

You must determine the cabling for the new controller modules and any new disk shelves to the existing FC switches.

#### About this task

This task must be performed at each MetroCluster site.

#### Steps

1. Use the procedure in [Fabric-attached MetroCluster installation and configuration](#) to create a cabling layout

for your switch type, using the port usage for an eight-node MetroCluster configuration.

The FC switch port usage must match the usage described in the procedure so that the Reference Configuration Files (RCFs) can be used.



If your environment cannot be cabled in such a way that RCF files can be used, you must manually configure the system according to instructions found in [Fabric-attached MetroCluster installation and configuration](#). Do not use this procedure if the cabling cannot use RCF files.

## Racking the new equipment

You must rack the equipment for the new nodes.

### Steps

1. Use the procedure in [Fabric-attached MetroCluster installation and configuration](#) to rack the new storage systems, disk shelves, and FC-to-SAS bridges.

## Verifying the health of the MetroCluster configuration

You should check the health of the MetroCluster configuration to verify proper operation.

### Steps

1. Check that the MetroCluster is configured and in normal mode on each cluster:

```
metrocluster show
```

```
cluster_A::> metrocluster show
Cluster Entry Name State

Local: cluster_A Configuration state configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-disaster
Remote: cluster_B Configuration state configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-disaster
```

2. Check that mirroring is enabled on each node:

```
metrocluster node show
```

```

cluster_A::> metrocluster node show
DR Configuration DR
Group Cluster Node State Mirroring Mode

1 cluster_A
 node_A_1 configured enabled normal
 cluster_B
 node_B_1 configured enabled normal
2 entries were displayed.

```

3. Check that the MetroCluster components are healthy:

```
metrocluster check run
```

```

cluster_A::> metrocluster check run

Last Checked On: 10/1/2014 16:03:37

Component Result

nodes ok
lifs ok
config-replication ok
aggregates ok
4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command
or sub-commands in "metrocluster check" directory for detailed results.
To check if the nodes are ready to do a switchover or switchback
operation, run "metrocluster switchover -simulate" or "metrocluster
switchback -simulate", respectively.

```

4. Check that there are no health alerts:

```
system health alert show
```

5. Simulate a switchover operation:

- a. From any node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with **y** when prompted to continue into advanced mode and see the advanced mode prompt (\*>).

- b. Perform the switchover operation with the `-simulate` parameter:

```
metrocluster switchover -simulate
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

## Checking for MetroCluster configuration errors with Config Advisor

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

### About this task

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

### Steps

1. Go to the Config Advisor download page and download the tool.

[NetApp Downloads: Config Advisor](#)

2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

## Sending a custom AutoSupport message prior to adding nodes to the MetroCluster configuration

You should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

### About this task

This task must be performed on each MetroCluster site.

### Steps

1. Log in to the cluster at Site\_A.
2. Invoke an AutoSupport message indicating the start of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours
```

The `maintenance-window-in-hours` parameter specifies the length of the maintenance window and can be a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can issue the following command to indicating that the maintenance period has ended:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

3. Repeat this step on the partner site.

## Recable and zone a switch fabric for the new nodes

### Disconnecting the existing DR group from the fabric

You must disconnect the existing controller modules from the FC switches in the fabric.

#### About this task

This task must be performed at each MetroCluster site.

#### Steps

1. Disable the HBA ports that connect the existing controller modules to the switch fabric undergoing maintenance:

```
storage port disable -node node-name -port port-number
```

2. On the local FC switches, remove the cables from the ports for the existing controller module's HBA, FC-VI, and ATTO bridges.

You should label the cables for easy identification when you re-cable them. Only the ISL ports should remain cabled.

### Recable and reconfigure the switches

You must apply the RCF files to reconfigure your zoning to accommodate the new nodes.

If you cannot use the RCF files to configure the switches, you must configure the switches manually. See:

- [Configure the Brocade FC switches manually](#)
- [Configure the Cisco FC switches manually](#)

#### Steps

1. Locate the RCF files for your configuration.

You must use the RCF files for an eight-node configuration and that match your switch model.

2. Apply the RCF files, following the directions on the download page, adjusting the ISL settings as needed.
3. Ensure that the switch configuration is saved.
4. Reboot the FC switches.
5. Cable both the pre-existing and the new FC-to-SAS bridges to the FC switches, using the cabling layout you created previously.

The FC switch port usage must match the MetroCluster eight-node usage described in [Fabric-attached MetroCluster installation and configuration](#) so that the Reference Configuration Files (RCFs) can be used.

6. Verify that the ports are online by using the correct command for your switch.

| Switch vendor | Command    |
|---------------|------------|
| Brocade       | switchshow |

|       |                      |
|-------|----------------------|
| Cisco | show interface brief |
|-------|----------------------|

7. Use the procedure in [Fabric-attached MetroCluster installation and configuration](#) to cable the FC-VI ports from the existing and new controllers, using the cabling layout you created previously.

The FC switch port usage must match the MetroCluster eight-node usage described in [Fabric-attached MetroCluster installation and configuration](#) so that the Reference Configuration Files (RCFs) can be used.

8. From the existing nodes, verify that the FC-VI ports are online:

```
metrocluster interconnect adapter show
```

```
metrocluster interconnect mirror show
```

9. Cable the HBA ports from the current and the new controllers.
10. On the existing controller modules, e-enable the ports connected to the switch fabric undergoing maintenance:

```
storage port enable -node node-name -port port-ID
```

11. Start the new controllers and boot them into Maintenance mode:

```
boot_ontap maint
```

12. Verify that only storage that will be used by the new DR group is visible to the new controller modules.

None of the storage that is used by the other DR group should be visible.

13. Return to the beginning of this process to re-cable the second switch fabric.

## Configure ONTAP on the new controllers

### Clearing the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing configuration.

#### Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond *yes* to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond *yes* to the confirmation prompt.

## Assigning disk ownership in AFF systems

If you are using AFF systems in a configuration with mirrored aggregates and the nodes do not have the disks (SSDs) correctly assigned, you should assign half the disks on each shelf to one local node and the other half of the disks to its HA partner node. You should create a configuration in which each node has the same number of disks in its local and remote disk pools.

### About this task

The storage controllers must be in Maintenance mode.

This does not apply to configurations which have unmirrored aggregates, an active/passive configuration, or that have an unequal number of disks in local and remote pools.

This task is not required if disks were correctly assigned when received from the factory.



Pool 0 always contains the disks that are found at the same site as the storage system that owns them, while Pool 1 always contains the disks that are remote to the storage system that owns them.

### Steps

1. If you have not done so, boot each system into Maintenance mode.
2. Assign the disks to the nodes located at the first site (site A):

You should assign an equal number of disks to each pool.

- a. On the first node, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0:

```
disk assign -disk disk-name -p pool -n number-of-disks
```

If storage controller Controller\_A\_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0 -n 4

*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf2 -p 1 -n 4
```

- b. Repeat the process for the second node at the local site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1:

```
disk assign -disk disk-name -p pool
```

If storage controller Controller\_A\_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1 -n 4

*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1 -n 4
```

3. Assign the disks to the nodes located at the second site (site B):

You should assign an equal number of disks to each pool.

- a. On the first node at the remote site, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0:

```
disk assign -disk disk-name -p pool
```

If storage controller Controller\_B\_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0 -n 4

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1 -n 4
```

- b. Repeat the process for the second node at the remote site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1:

```
disk assign -disk disk-name -p pool
```

If storage controller Controller\_B\_2 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0 -n 4

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1 -n 4
```

4. Confirm the disk assignments:

```
storage show disk
```

5. Exit Maintenance mode:

```
halt
```

6. Display the boot menu:

```
boot_ontap menu
```

7. On each node, select option 4 to initialize all disks.

### Assigning disk ownership in non-AFF systems

If the MetroCluster nodes do not have the disks correctly assigned, or if you are using DS460C disk shelves in your configuration, you must assign disks to each of the nodes in the MetroCluster configuration on a shelf-by-shelf basis. You will create a configuration in which each node has the same number of disks in its local and remote disk pools.

#### About this task

The storage controllers must be in Maintenance mode.

If your configuration does not include DS460C disk shelves, this task is not required if disks were correctly assigned when received from the factory.



Pool 0 always contains the disks that are found at the same site as the storage system that owns them.

Pool 1 always contains the disks that are remote to the storage system that owns them.

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

| Assign these disks in the drawer... | To this node and pool...              |
|-------------------------------------|---------------------------------------|
| 0 - 2                               | Local node's pool 0                   |
| 3 - 5                               | HA partner node's pool 0              |
| 6 - 8                               | DR partner of the local node's pool 1 |

This disk assignment pattern ensures that an aggregate is minimally affected in case a drawer goes offline.

### Steps

1. If you have not done so, boot each system into Maintenance mode.
2. Assign the disk shelves to the nodes located at the first site (site A):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

If storage controller Controller\_A\_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0

*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1
*> disk assign -shelf FC_switch_B_1:1-4.shelf2 -p 1
```

- b. Repeat the process for the second node at the local site, systematically assigning the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

If storage controller Controller\_A\_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1

*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1
```

3. Assign the disk shelves to the nodes located at the second site (site B):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node at the remote site, systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf local-switch-nameshelf-name -p pool
```

If storage controller Controller\_B\_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1
```

- b. Repeat the process for the second node at the remote site, systematically assigning its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf shelf-name -p pool
```

If storage controller Controller\_B\_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1
```

4. Confirm the shelf assignments:

```
storage show shelf
```

5. Exit Maintenance mode:

```
halt
```

6. Display the boot menu:

```
boot_ontap menu
```

7. On each node, select option 4 to initialize all disks.

### Verifying the ha-config state of components

In a MetroCluster configuration, the ha-config state of the controller module and chassis components must be set to **mcc** so they boot up properly.

#### About this task

- The system must be in Maintenance mode.
- This task must be performed on each new controller module.

#### Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be "mcc".

2. If the displayed system state of the controller is not correct, set the HA state for the controller module:

```
ha-config modify controller mcc
```

3. If the displayed system state of the chassis is not correct, set the HA state for the chassis:

```
ha-config modify chassis mcc
```

4. Repeat these steps on the other replacement node.

### **Booting the new controllers and joining them to the cluster**

To join the new controllers to the cluster, you must boot each new controller module and use the ONTAP cluster setup wizard to identify the cluster will join.

#### **Before you begin**

You must have cabled the MetroCluster configuration.

You must not have configured the Service Processor prior to performing this task.

#### **About this task**

This task must be performed on each of the new controllers at both clusters in the MetroCluster configuration.

#### **Steps**

1. If you have not already done so, power up each node and let them boot completely.

If the system is in Maintenance mode, issue the `halt` command to exit Maintenance mode, and then issue the following command from the LOADER prompt:

```
boot_ontap
```

The controller module enters the node setup wizard.

The output should be similar to the following:

```
Welcome to node setup
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the setup wizard.
```

```
Any changes you made before quitting will be saved.
```

```
To accept a default or omit a question, do not enter a value.
```

```
.
. .
.
```

2. Enable the AutoSupport tool by following the directions provided by the system.
3. Respond to the prompts to configure the node management interface.

The prompts are similar to the following:

```
Enter the node management interface port: [e0M]:
Enter the node management interface IP address: 10.228.160.229
Enter the node management interface netmask: 225.225.252.0
Enter the node management interface default gateway: 10.228.160.1
```

4. Confirm that nodes are configured in high-availability mode:

```
storage failover show -fields mode
```

If not, you must issue the following command on each node, and then reboot the node:

```
storage failover modify -mode ha -node localhost
```

This command configures high availability mode but does not enable storage failover. Storage failover is automatically enabled when you issue the `metrocluster configure` command later in the configuration process.

5. Confirm that you have four ports configured as cluster interconnects:

```
network port show
```

The following example shows output for two controllers in `cluster_A`. If it is a two-node MetroCluster configuration, the output shows only one node.

```
cluster_A::> network port show
```

```
(Mbps)
Node Port IPspace Broadcast Domain Link MTU Admin/Oper

node_A_1
 **e0a Cluster Cluster up 1500
auto/1000
 e0b Cluster Cluster up 1500
auto/1000**
 e0c Default Default up 1500 auto/1000
 e0d Default Default up 1500 auto/1000
 e0e Default Default up 1500 auto/1000
 e0f Default Default up 1500 auto/1000
 e0g Default Default up 1500 auto/1000
node_A_2
 **e0a Cluster Cluster up 1500
auto/1000
 e0b Cluster Cluster up 1500
auto/1000**
 e0c Default Default up 1500 auto/1000
 e0d Default Default up 1500 auto/1000
 e0e Default Default up 1500 auto/1000
 e0f Default Default up 1500 auto/1000
 e0g Default Default up 1500 auto/1000
14 entries were displayed.
```

6. Because you are using the CLI to set up the cluster, exit the Node Setup wizard:

```
exit
```

7. Log in to the admin account by using the `admin` user name.

8. Start the Cluster Setup wizard, and then join the existing cluster:

```
cluster setup
```

```
::> cluster setup
```

Welcome to the cluster setup wizard.

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,  
"back" - if you want to change previously answered questions, and  
"exit" or "quit" - if you want to quit the cluster setup wizard.  
Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".  
To accept a default or omit a question, do not enter a value.

Do you want to create a new cluster or join an existing cluster?  
{create, join}:`join`

9. After you complete the **Cluster Setup** wizard and it exits, verify that the cluster is active and the node is healthy:

```
cluster show
```

The following example shows a cluster in which the first node (cluster1-01) is healthy and eligible to participate:

```
cluster_A::> cluster show
Node Health Eligibility

node_A_1 true true
node_A_2 true true
node_A_3 true true
```

If it becomes necessary to change any of the settings you entered for the admin SVM or node SVM, you can access the **Cluster Setup** wizard by using the `cluster setup` command.

## Configure the clusters into a MetroCluster configuration

### Configure intercluster LIFs

Learn how to configure intercluster LIFs on dedicated and shared ports.

### **Configure intercluster LIFs on dedicated ports**

You can configure intercluster LIFs on dedicated ports to increase the available bandwidth for replication traffic.

#### **Steps**

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```

cluster01::> network port show

```

|              |      |         |                  |      |      | Speed |
|--------------|------|---------|------------------|------|------|-------|
| (Mbps)       |      |         |                  |      |      |       |
| Node         | Port | IPspace | Broadcast Domain | Link | MTU  |       |
| Admin/Oper   |      |         |                  |      |      |       |
| -----        |      |         |                  |      |      | ----- |
| cluster01-01 |      |         |                  |      |      |       |
|              | e0a  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0b  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0c  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0d  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0e  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0f  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
| cluster01-02 |      |         |                  |      |      |       |
|              | e0a  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0b  | Cluster | Cluster          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0c  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0d  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0e  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |
|              | e0f  | Default | Default          | up   | 1500 |       |
| auto/1000    |      |         |                  |      |      |       |

2. Determine which ports are available to dedicate to intercluster communication:

```
network interface show -fields home-port,curr-port
```

For complete command syntax, see the man page.

The following example shows that ports "e0e" and "e0f" have not been assigned LIFs:

```

cluster01::> network interface show -fields home-port,curr-port
vserver lif home-port curr-port

Cluster cluster01-01_clus1 e0a e0a
Cluster cluster01-01_clus2 e0b e0b
Cluster cluster01-02_clus1 e0a e0a
Cluster cluster01-02_clus2 e0b e0b
cluster01
 cluster_mgmt e0c e0c
cluster01
 cluster01-01_mgmt1 e0c e0c
cluster01
 cluster01-02_mgmt1 e0c e0c

```

### 3. Create a failover group for the dedicated ports:

```

network interface failover-groups create -vserver <system_SVM> -failover
-group <failover_group> -targets <physical_or_logical_ports>

```

The following example assigns ports "e0e" and "e0f" to the failover group "intercluster01" on the system SVM "cluster01":

```

cluster01::> network interface failover-groups create -vserver
cluster01 -failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f

```

### 4. Verify that the failover group was created:

```

network interface failover-groups show

```

For complete command syntax, see the man page.

```

cluster01::> network interface failover-groups show
 Failover
Vserver Group Targets

Cluster
 Cluster
 cluster01-01:e0a, cluster01-
01:e0b,
 cluster01-02:e0a, cluster01-02:e0b
cluster01
 Default
 cluster01-01:e0c, cluster01-
01:e0d,
 cluster01-02:e0c, cluster01-
02:e0d,
 cluster01-01:e0e, cluster01-01:e0f
 cluster01-02:e0e, cluster01-02:e0f
 intercluster01
 cluster01-01:e0e, cluster01-01:e0f
 cluster01-02:e0e, cluster01-02:e0f

```

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

| ONTAP version   | Command                                                                                                                                                                                                        |
|-----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 9.6 and later   | network interface create -vserver <system_SVM> -lif <LIF_name> -service-policy default-intercluster -home-node <node> -home-port <port> -address <port_IP> -netmask <netmask> -failover-group <failover_group> |
| 9.5 and earlier | network interface create -vserver system_SVM -lif <LIF_name> -role intercluster -home-node <node> -home-port <port> -address <port_IP> -netmask <netmask> -failover-group <failover_group>                     |

For complete command syntax, see the man page.

The following example creates intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" in the failover group "intercluster01":

```

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0e
-address 192.168.1.201
-netmask 255.255.255.0 -failover-group intercluster01

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0e
-address 192.168.1.202
-netmask 255.255.255.0 -failover-group intercluster01

```

6. Verify that the intercluster LIFs were created:

**In ONTAP 9.6 and later:**

```
network interface show -service-policy default-intercluster
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster
```

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-
intercluster

```

| Current Is | Logical         | Status     | Network          | Current      |
|------------|-----------------|------------|------------------|--------------|
| Vserver    | Interface       | Admin/Oper | Address/Mask     | Node         |
| Port       | Home            |            |                  |              |
| -----      | -----           | -----      | -----            | -----        |
| -----      | -----           | -----      | -----            | -----        |
| cluster01  | cluster01_icl01 | up/up      | 192.168.1.201/24 | cluster01-01 |
| e0e        | true            |            |                  |              |
|            | cluster01_icl02 | up/up      | 192.168.1.202/24 | cluster01-02 |
| e0f        | true            |            |                  |              |

7. Verify that the intercluster LIFs are redundant:

**In ONTAP 9.6 and later:**

```
network interface show -service-policy default-intercluster -failover
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" on the SVM "e0e" port will fail over to the "e0f" port.

```
cluster01::> network interface show -service-policy default-
intercluster -failover
 Logical Home Failover
Failover
Vserver Interface Node:Port Policy Group
----- -----
cluster01
 cluster01_icl01 cluster01-01:e0e local-only
intercluster01
 Failover Targets: cluster01-01:e0e,
 cluster01-01:e0f
 cluster01_icl02 cluster01-02:e0e local-only
intercluster01
 Failover Targets: cluster01-02:e0e,
 cluster01-02:e0f
```

### Configure intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network to reduce the number of ports you need for intercluster networking.

#### Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```
cluster01::> network port show
```

|              |       |         |                  |       |       | Speed |
|--------------|-------|---------|------------------|-------|-------|-------|
| (Mbps)       |       |         |                  |       |       |       |
| Node         | Port  | IPspace | Broadcast Domain | Link  | MTU   |       |
| Admin/Oper   |       |         |                  |       |       |       |
| -----        | ----- | -----   | -----            | ----- | ----- | ----- |
| cluster01-01 |       |         |                  |       |       |       |
|              | e0a   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0b   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0c   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0d   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
| cluster01-02 |       |         |                  |       |       |       |
|              | e0a   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0b   | Cluster | Cluster          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0c   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |
|              | e0d   | Default | Default          | up    | 1500  |       |
| auto/1000    |       |         |                  |       |       |       |

## 2. Create intercluster LIFs on the system SVM:

### In ONTAP 9.6 and later:

```
network interface create -vserver <system_SVM> -lif <LIF_name> -service
-policy default-intercluster -home-node <node> -home-port <port> -address
<port_IP> -netmask <netmask>
```

### In ONTAP 9.5 and earlier:

```
network interface create -vserver <system_SVM> -lif <LIF_name> -role
intercluster -home-node <node> -home-port <port> -address <port_IP>
-netmask <netmask>
```

For complete command syntax, see the man page.

The following example creates intercluster LIFs `cluster01_icl01` and `cluster01_icl02`:

```

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl01 -service-
policy default-intercluster -home-node cluster01-01 -home-port e0c
-address 192.168.1.201
-netmask 255.255.255.0

cluster01::> network interface create -vserver cluster01 -lif
cluster01_icl02 -service-
policy default-intercluster -home-node cluster01-02 -home-port e0c
-address 192.168.1.202
-netmask 255.255.255.0

```

### 3. Verify that the intercluster LIFs were created:

#### In ONTAP 9.6 and later:

```
network interface show -service-policy default-intercluster
```

#### In ONTAP 9.5 and earlier:

```
network interface show -role intercluster
```

For complete command syntax, see the man page.

```

cluster01::> network interface show -service-policy default-
intercluster

 Logical Status Network Current
Current Is
Vserver Interface Admin/Oper Address/Mask Node
Port Home

cluster01
 cluster01_icl01
 up/up 192.168.1.201/24 cluster01-01
e0c true
 cluster01_icl02
 up/up 192.168.1.202/24 cluster01-02
e0c true

```

### 4. Verify that the intercluster LIFs are redundant:

#### In ONTAP 9.6 and later:

```
network interface show -service-policy default-intercluster -failover
```

**In ONTAP 9.5 and earlier:**

```
network interface show -role intercluster -failover
```

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" on the "e0c" port will fail over to the "e0d" port.

```
cluster01::> network interface show -service-policy default-
intercluster -failover
 Logical Home Failover
Failover
Vserver Interface Node:Port Policy Group
----- -----
cluster01
 cluster01_icl01 cluster01-01:e0c local-only
192.168.1.201/24
 Failover Targets: cluster01-01:e0c,
 cluster01-01:e0d
 cluster01_icl02 cluster01-02:e0c local-only
192.168.1.201/24
 Failover Targets: cluster01-02:e0c,
 cluster01-02:e0d
```

### Mirroring the root aggregates

You must mirror the root aggregates to provide data protection.

By default, the root aggregate is created as RAID-DP type aggregate. You can change the root aggregate from RAID-DP to RAID4 type aggregate. The following command modifies the root aggregate for RAID4 type aggregate:

```
storage aggregate modify -aggregate aggr_name -raidtype raid4
```



On non-ADP systems, the RAID type of the aggregate can be modified from the default RAID-DP to RAID4 before or after the aggregate is mirrored.

### Steps

1. Mirror the root aggregate:

```
storage aggregate mirror aggr_name
```

The following command mirrors the root aggregate for controller\_A\_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote MetroCluster site.

2. Repeat the previous step for each node in the MetroCluster configuration.

### Implementing the MetroCluster configuration

You must run the `metrocluster configure -refresh true` command to start data protection on the nodes that you have added to a MetroCluster configuration.

#### About this task

You issue the `metrocluster configure -refresh true` command once, on one of the newly added nodes, to refresh the MetroCluster configuration. You do not need to issue the command on each of the sites or nodes.

The `metrocluster configure -refresh true` command automatically pairs the two nodes with the lowest system IDs in each of the two clusters as disaster recovery (DR) partners. In a four-node MetroCluster configuration, there are two DR partner pairs. The second DR pair is created from the two nodes with higher system IDs.

#### Steps

1. Refresh the MetroCluster configuration:

- a. Enter advanced privilege mode:

```
set -privilege advanced
```

- b. Refresh the MetroCluster configuration on one of the new nodes:

```
metrocluster configure -refresh true
```

The following example shows the MetroCluster configuration refreshed on both DR groups:

```
controller_A_2::*> metrocluster configure -refresh true
[Job 726] Job succeeded: Configure is successful.
```

```
controller_A_4::*> metrocluster configure -refresh true
[Job 740] Job succeeded: Configure is successful.
```

- c. Return to admin privilege mode:

```
set -privilege admin
```

## 2. Verify the networking status on site A:

```
network port show
```

The following example shows the network port usage on a four-node MetroCluster configuration:

```
cluster_A::> network port show
```

| Node           | Port | IPspace | Broadcast Domain | Link | MTU  | Speed (Mbps)<br>Admin/Oper |
|----------------|------|---------|------------------|------|------|----------------------------|
| controller_A_1 |      |         |                  |      |      |                            |
|                | e0a  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0b  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0c  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0d  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0e  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0f  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0g  | Default | Default          | up   | 1500 | auto/1000                  |
| controller_A_2 |      |         |                  |      |      |                            |
|                | e0a  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0b  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0c  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0d  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0e  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0f  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0g  | Default | Default          | up   | 1500 | auto/1000                  |

14 entries were displayed.

## 3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration:

### a. Verify the configuration from site A:

```
metrocluster show
```

```
cluster_A::> metrocluster show
```

Configuration: IP fabric

| Cluster           | Entry Name          | State      |
|-------------------|---------------------|------------|
| Local: cluster_A  | Configuration state | configured |
|                   | Mode                | normal     |
| Remote: cluster_B | Configuration state | configured |
|                   | Mode                | normal     |

b. Verify the configuration from site B:

```
metrocluster show
```

```
cluster_B::> metrocluster show
```

```
Configuration: IP fabric
```

| Cluster           | Entry Name          | State      |
|-------------------|---------------------|------------|
| Local: cluster_B  | Configuration state | configured |
|                   | Mode                | normal     |
| Remote: cluster_A | Configuration state | configured |
|                   | Mode                | normal     |

### Create a mirrored data aggregate on each MetroCluster FC node

You must create a mirrored data aggregate on each node in the DR group.

#### About this task

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can make sure that the correct drive type is selected.
- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.

In systems using ADP, aggregates are created using partitions in which each drive is partitioned in to P1, P2 and P3 partitions.

- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.

#### Disk and aggregate management

- Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.



It's recommended you maintain at least 20% free space for mirrored aggregates for optimal storage performance and availability. Although the recommendation is 10% for non-mirrored aggregates, the additional 10% of space can be used by the filesystem to absorb incremental changes. Incremental changes increase space utilization for mirrored aggregates due to ONTAP's copy-on-write Snapshot-based architecture. Failure to adhere to these best practices might have a negative impact on performance.

#### Steps

1. Display a list of available spares:

```
storage disk show -spare -owner <node_name>
```

## 2. Create the aggregate:

```
storage aggregate create -mirror true
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To verify that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include



In the minimum supported configuration, in which a limited number of drives are available, you must use the `force-small-aggregate` option to allow the creation of a three disk RAID-DP aggregate.

- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create` man page.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10
-node node_A_1 -mirror true
[Job 15] Job is queued: Create aggr1_node_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

## 3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate <aggregate-name>
```

### Configuring FC-to-SAS bridges for health monitoring

Learn how to configure the FC-to-SAS bridges for health monitoring.

#### About this task

- Third-party SNMP monitoring tools are not supported for FibreBridge bridges.

- Beginning with ONTAP 9.8, FC-to-SAS bridges are monitored via in-band connections by default, and additional configuration is not required.



Beginning with ONTAP 9.8, the `storage bridge` command is replaced with `system bridge`. The following steps show the `storage bridge` command, but if you are running ONTAP 9.8 or later, the `system bridge` command is preferred.

### Step

- From the ONTAP cluster prompt, add the bridge to health monitoring:
  - Add the bridge, using the command for your version of ONTAP:

| ONTAP version   | Command                                                                                       |
|-----------------|-----------------------------------------------------------------------------------------------|
| 9.5 and later   | <code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code> |
| 9.4 and earlier | <code>storage bridge add -address <i>bridge-ip-address</i> -name <i>bridge-name</i></code>    |

- Verify that the bridge has been added and is properly configured:

```
storage bridge show
```

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the "Status" column is "ok", and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```
controller_A_1::> storage bridge show

Bridge Symbolic Name Is Monitored Monitor Status
Vendor Model Bridge WWN

ATTO_10.10.20.10 atto01 true ok Atto
FibreBridge 7500N 20000010867038c0
ATTO_10.10.20.11 atto02 true ok Atto
FibreBridge 7500N 20000010867033c0
ATTO_10.10.20.12 atto03 true ok Atto
FibreBridge 7500N 20000010867030c0
ATTO_10.10.20.13 atto04 true ok Atto
FibreBridge 7500N 2000001086703b80

4 entries were displayed

controller_A_1::>
```

## Moving a metadata volume in MetroCluster configurations

You can move a metadata volume from one aggregate to another aggregate in a MetroCluster configuration. You might want to move a metadata volume when the source aggregate is decommissioned or unmirrored, or for other reasons that make the aggregate ineligible.

### About this task

- You must have cluster administrator privileges to perform this task.
- The target aggregate must be mirrored and should not be in the degraded state.
- The available space in the target aggregate must be larger than the metadata volume that you are moving.

### Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Identify the metadata volume that should be moved:

```
volume show MDV_CRS*
```

```

Cluster_A::*> volume show MDV_CRS*
Vserver Volume Aggregate State Type Size
Available Used%

Cluster_A
 MDV_CRS_14c00d4ac9f311e7922800a0984395f1_A
 Node_A_1_aggr1
 online RW 10GB
9.50GB 5%
Cluster_A
 MDV_CRS_14c00d4ac9f311e7922800a0984395f1_B
 Node_A_2_aggr1
 online RW 10GB
9.50GB 5%
Cluster_A
 MDV_CRS_15035e66c9f311e7902700a098439625_A
 Node_B_1_aggr1
 - RW -
- -
Cluster_A
 MDV_CRS_15035e66c9f311e7902700a098439625_B
 Node_B_2_aggr1
 - RW -
- -
4 entries were displayed.

Cluster_A::>

```

### 3. Identify an eligible target aggregate:

```
metrocluster check config-replication show-aggregate-eligibility
```

The following command identifies the aggregates in cluster\_A that are eligible to host metadata volumes:

```
Cluster_A::*> metrocluster check config-replication show-aggregate-eligibility
```

```
Aggregate Hosted Config Replication Vols Host Addl Vols Comments

Node_A_1_aggr0 - false Root Aggregate
Node_A_2_aggr0 - false Root Aggregate
Node_A_1_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_A true -
Node_A_2_aggr1 MDV_CRS_1bc7134a5ddf11e3b63f123478563412_B true -
Node_A_1_aggr2 - true
Node_A_2_aggr2 - true
Node_A_1_Aggr3 - false Unable to determine available space of aggregate
Node_A_1_aggr5 - false Unable to determine mirror configuration
Node_A_2_aggr6 - false Mirror configuration does not match requirement
Node_B_1_aggr4 - false NonLocal Aggregate
```



In the previous example, Node\_A\_1\_aggr2 and Node\_A\_2\_aggr2 are eligible.

#### 4. Start the volume move operation:

```
volume move start -vserver svm_name -volume metadata_volume_name -destination
-aggregate destination_aggregate_name*
```

The following command moves metadata volume "MDV\_CRS\_14c00d4ac9f311e7922800a0984395f1" from "aggregate Node\_A\_1\_aggr1" to "aggregate Node\_A\_1\_aggr2":

```
Cluster_A::*> volume move start -vserver svm_cluster_A -volume
MDV_CRS_14c00d4ac9f311e7922800a0984395f1
-destination-aggregate aggr_cluster_A_02_01

Warning: You are about to modify the system volume
"MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A". This may cause
severe
performance or stability problems. Do not proceed unless
directed to
do so by support. Do you want to proceed? {y|n}: y
[Job 109] Job is queued: Move
"MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" in Vserver
"svm_cluster_A" to aggregate "aggr_cluster_A_02_01".
Use the "volume move show -vserver svm_cluster_A -volume
MDV_CRS_9da04864ca6011e7b82e0050568be9fe_A" command to view the status
of this operation.
```

#### 5. Verify the state of the volume move operation:

```
volume move show -volume vol_constituent_name
```

6. Return to the admin privilege level:

```
set -privilege admin
```

### Checking the MetroCluster configuration

You can check that the components and relationships in the MetroCluster configuration are working correctly. You should do a check after initial configuration and after making any changes to the MetroCluster configuration. You should also do a check before a negotiated (planned) switchover or a switchback operation.

#### About this task

If the `metrocluster check run` command is issued twice within a short time on either or both clusters, a conflict can occur and the command might not collect all data. Subsequent `metrocluster check show` commands do not show the expected output.

#### Steps

1. Check the configuration:

```
metrocluster check run
```

The command runs as a background job and might not be completed immediately.

```
cluster_A::> metrocluster check run
The operation has been started and is running in the background. Wait
for
it to complete and run "metrocluster check show" to view the results. To
check the status of the running metrocluster check operation, use the
command,
"metrocluster operation history show -job-id 2245"
```

```
cluster_A::> metrocluster check show
```

| Component          | Result |
|--------------------|--------|
| nodes              | ok     |
| lifs               | ok     |
| config-replication | ok     |
| aggregates         | ok     |
| clusters           | ok     |
| connections        | ok     |
| volumes            | ok     |

7 entries were displayed.

2. Display more detailed results from the most recent `metrocluster check run` command:

```
metrocluster check aggregate show
```

```
metrocluster check cluster show
```

```
metrocluster check config-replication show
```

```
metrocluster check lif show
```

```
metrocluster check node show
```

The `metrocluster check show` commands show the results of the most recent `metrocluster check run` command. You should always run the `metrocluster check run` command prior to using the `metrocluster check show` commands so that the information displayed is current.

The following example shows the `metrocluster check aggregate show` command output for a healthy four-node MetroCluster configuration:

```
cluster_A::> metrocluster check aggregate show

Last Checked On: 8/5/2014 00:42:58

Node Aggregate Check
Result

controller_A_1 controller_A_1_aggr0
ok
ok
ok
controller_A_1_aggr1
ok
ok
ok
controller_A_1_aggr2
ok
ok
ok
```

```

controller_A_2 controller_A_2_aggr0
ok
ok
ok
ok
controller_A_2_aggr1
ok
ok
ok
ok
controller_A_2_aggr2
ok
ok
ok
ok
18 entries were displayed.

```

The following example shows the `metrocluster check cluster show` command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

Last Checked On: 9/13/2017 20:47:04

| Cluster             | Check                       | Result         |
|---------------------|-----------------------------|----------------|
| mccint-fas9000-0102 | negotiated-switchover-ready | not-applicable |
|                     | switchback-ready            | not-applicable |
|                     | job-schedules               | ok             |
|                     | licenses                    | ok             |
|                     | periodic-check-enabled      | ok             |
| mccint-fas9000-0304 | negotiated-switchover-ready | not-applicable |
|                     | switchback-ready            | not-applicable |
|                     | job-schedules               | ok             |
|                     | licenses                    | ok             |
|                     | periodic-check-enabled      | ok             |

10 entries were displayed.

## Checking for MetroCluster configuration errors with Config Advisor

You can go to the NetApp Support Site and download the Config Advisor tool to check for common configuration errors.

### About this task

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.



Support for Config Advisor is limited, and available only online.

### Steps

1. Go to the Config Advisor download page and download the tool.

[NetApp Downloads: Config Advisor](#)

2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

## Sending a custom AutoSupport message after to adding nodes to the MetroCluster configuration

You should issue an AutoSupport message to notify NetApp technical support that maintenance is complete.

### About this task

This task must be performed on each MetroCluster site.

### Steps

1. Log in to the cluster at Site\_A.
2. Invoke an AutoSupport message indicating the end of the maintenance:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

3. Repeat this step on the partner site.

## Verifying switchover, healing, and switchback

You should verify the switchover, healing, and switchback operations of the MetroCluster configuration.

### Steps

1. Use the procedures for negotiated switchover, healing, and switchback in [MetroCluster management and disaster recovery](#).

## Expand a MetroCluster IP configuration

Depending on your ONTAP version, you can expand your MetroCluster IP configuration by adding four new nodes as a new DR group.

Beginning with ONTAP 9.13.1, you can temporarily expand an eight-node MetroCluster configuration to refresh the controllers and storage. See [Refreshing a four-node or an eight-node MetroCluster IP configuration \(ONTAP 9.8 and later\)](#) for more information.

Beginning with ONTAP 9.9.1, you can add four new nodes to the MetroCluster IP configuration as a second DR group. This creates an eight-node MetroCluster configuration.

### Important information if you are adding an older platform model

The following guidance is for an uncommon scenario where you need to add an older platform model (platforms released before ONTAP 9.15.1) to an existing MetroCluster configuration that contains a newer platform model (platforms released in ONTAP 9.15.1 or later).

If your existing MetroCluster configuration contains a platform that uses **shared cluster/HA ports** (platforms released in ONTAP 9.15.1 or later), you cannot add a platform that uses **shared MetroCluster/HA ports** (platforms released before ONTAP 9.15.1) without upgrading all nodes in the configuration to ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later.



Adding an older platform model that uses **shared/MetroCluster HA ports** to a MetroCluster containing a newer platform model that uses **shared cluster/HA ports** is an uncommon scenario and most combinations are not affected.

Use the following table to verify whether your combination is affected. If your existing platform is listed in the first column, and the platform you want to add to the configuration is listed in the second column, all nodes in the configuration must be running ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later to add the new DR group.

| If your existing MetroCluster contains..                                                                                                                                                                                                                                 |                                                                                                                                                 | And the platform you're adding is...                                                                                                                                                                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                        | Then...                                                                                                                                                                             |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| An AFF system using <b>shared cluster/HA ports</b> : <ul style="list-style-type: none"> <li>• AFF A20</li> <li>• AFF A30</li> <li>• AFF C30</li> <li>• AFF A50</li> <li>• AFF C60</li> <li>• AFF C80</li> <li>• AFF A70</li> <li>• AFF A90</li> <li>• AFF A1K</li> </ul> | A FAS system using <b>shared cluster/HA ports</b> : <ul style="list-style-type: none"> <li>• FAS50</li> <li>• FAS70</li> <li>• FAS90</li> </ul> | An AFF system using <b>shared MetroCluster/HA ports</b> : <ul style="list-style-type: none"> <li>• AFF A150, ASA A150</li> <li>• AFF A220</li> <li>• AFF C250, ASA C250</li> <li>• AFF A250, ASA A250</li> <li>• AFF A300</li> <li>• AFF A320</li> <li>• AFF C400, ASA C400</li> <li>• AFF A400, ASA A400</li> <li>• AFF A700</li> <li>• AFF C800, ASA C800</li> <li>• AFF A800, ASA A800</li> <li>• AFF A900, ASA A900</li> </ul> | A FAS system using <b>shared MetroCluster/HA ports</b> : <ul style="list-style-type: none"> <li>• FAS2750</li> <li>• FAS500f</li> <li>• FAS8200</li> <li>• FAS8300</li> <li>• FAS8700</li> <li>• FAS9000</li> <li>• FAS9500</li> </ul> | Before you add the new platform to your existing MetroCluster configuration, upgrade all nodes in the existing and new configuration to ONTAP 9.15.1P11 or ONTAP 9.16.1P4 or later. |

### Before you begin

- The old and new nodes must be running the same version of ONTAP.
- This procedure describes the steps required to add one four-node DR group to an existing MetroCluster IP configuration. If you are refreshing an eight-node configuration, you must repeat the entire procedure for each DR group, adding one at a time.
- Verify that the old and new platform models are supported for platform mixing.

[NetApp Hardware Universe](#)

- Verify that the old and new platform models are both supported by the IP switches.

[NetApp Hardware Universe](#)

- If you are [refreshing a four-node or an eight-node MetroCluster IP configuration](#), the new nodes must have enough storage to accommodate the data of the old nodes, along with adequate disks for root aggregates and spare disks.
- Verify that you have a default broadcast domain created on the old nodes.

When you add new nodes to an existing cluster without a default broadcast domain, node management

LIFs are created for the new nodes using universal unique identifiers (UUIDs) instead of the expected names. For more information, see the Knowledge Base article [Node management LIFs on newly-added nodes generated with UUID names](#).

## Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Example naming in this procedure

This procedure uses example names throughout to identify the DR groups, nodes, and switches involved.

| DR groups      | cluster_A at site_A                                                                   | cluster_B at site_B                                                                   |
|----------------|---------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| dr_group_1-old | <ul style="list-style-type: none"><li>• node_A_1-old</li><li>• node_A_2-old</li></ul> | <ul style="list-style-type: none"><li>• node_B_1-old</li><li>• node_B_2-old</li></ul> |
| dr_group_2-new | <ul style="list-style-type: none"><li>• node_A_3-new</li><li>• node_A_4-new</li></ul> | <ul style="list-style-type: none"><li>• node_B_3-new</li><li>• node_B_4-new</li></ul> |

## Supported platform combinations when adding a second DR group

The following tables shows the supported platform combinations for eight-node MetroCluster IP configurations.



- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, if you have an eight-node configuration, all eight nodes must be running the same ONTAP version. Refer to the [Hardware universe](#) for the minimum supported ONTAP version for your combination.
- The combinations in this table apply only to regular or permanent eight-node configurations.
- The platform combinations shown in this table **do not** apply if you are using the transition or refresh procedures.
- All nodes in one DR group must be of the same type and configuration.

## Supported AFF and FAS MetroCluster IP expansion combinations

The following tables show the supported platform combinations for expanding an AFF or FAS system in a MetroCluster IP configuration. The tables are split into two groups:

- **Group 1** shows combinations for AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.
- **Group 2** shows combinations for AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

The following notes apply to both groups:

- Note 1: ONTAP 9.9.1 or later (or the minimum ONTAP version supported on the platform) is required for these combinations.
- Note 2: ONTAP 9.13.1 or later (or the minimum ONTAP version supported on the platform) is required for these combinations.

### AFF and FAS combinations group 1

Review the expansion combinations for AFF A150, AFF A20, FAS2750, AFF A220, FAS500f, AFF C250, AFF A250, FAS50, AFF C30, AFF A30, FAS8200, AFF A300, AFF A320, FAS8300, AFF C400, AFF A400, and FAS8700 systems.

| AFF and FAS              |                                 | Eight-node DR group 2 |         |                     |                                 |       |                    |                     |          |                                 |         |
|--------------------------|---------------------------------|-----------------------|---------|---------------------|---------------------------------|-------|--------------------|---------------------|----------|---------------------------------|---------|
|                          |                                 | AFF A150              | AFF A20 | FAS2750<br>AFF A220 | FAS500f<br>AFF C250<br>AFF A250 | FAS50 | AFF C30<br>AFF A30 | FAS8200<br>AFF A300 | AFF A320 | FAS8300<br>AFF C400<br>AFF A400 | FAS8700 |
| Eight-node<br>DR group 1 | AFF A150                        | Note 2                |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | AFF A20                         | Note 2                |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS2750<br>AFF A220             | Note 2                |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS500f<br>AFF C250<br>AFF A250 | Note 2                |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS50                           | Note 2                |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | AFF C30<br>AFF A30              | Note 2                |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS8200<br>AFF A300             |                       |         |                     |                                 |       |                    | Note 1              |          |                                 |         |
|                          | AFF A320                        |                       |         |                     |                                 |       |                    | Note 1              |          |                                 |         |
|                          | FAS8300<br>AFF C400<br>AFF A400 |                       |         |                     |                                 |       |                    | Note 1              | Note 1   |                                 |         |
|                          | FAS8700                         |                       |         |                     |                                 |       |                    |                     | Note 1   |                                 |         |
|                          | AFF C60                         |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | AFF A50                         |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS70                           |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS9000<br>AFF A700<br>AFF A70  |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | AFF C800<br>AFF A800            |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS9500<br>AFF A900<br>AFF C80  |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | FAS90                           |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | AFF A90                         |                       |         |                     |                                 |       |                    |                     |          |                                 |         |
|                          | AFF A1K                         |                       |         |                     |                                 |       |                    |                     |          |                                 |         |

### AFF and FAS combinations group 2

Review the expansion combinations for AFF C60, AFF A50, FAS70, FAS9000, AFF A700, AFF A70, AFF C800, AFF A800, FAS9500, AFF A900, AFF C80, FAS90, AFF A90, and AFF A1K systems.

| AFF and FAS              |          | Eight-node DR group 2 |         |        |                     |         |                      |                     |         |                  |         |
|--------------------------|----------|-----------------------|---------|--------|---------------------|---------|----------------------|---------------------|---------|------------------|---------|
|                          |          | AFF C60               | AFF A50 | FAS70  | FAS9000<br>AFF A700 | AFF A70 | AFF C800<br>AFF A800 | FAS9500<br>AFF A900 | AFF C80 | FAS90<br>AFF A90 | AFF A1K |
| Eight-node<br>DR group 1 | AFF A150 |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A20  |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS2750  |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A220 |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS500f  |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF C250 |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A250 |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS50    |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF C30  |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A30  |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS8200  |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A300 |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A320 |                       |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS8300  | Note 1                |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF C400 | Note 1                |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A400 | Note 1                |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS8700  | Note 1                |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF C60  | Note 1                |         |        |                     |         |                      |                     |         |                  |         |
|                          | AFF A50  | Note 1                |         |        |                     |         |                      |                     |         |                  |         |
|                          | FAS70    | Note 1                |         |        | Note 1              |         |                      |                     |         |                  |         |
| FAS9000                  |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF A700                 |          |                       | Note 1  | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF A70                  |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF C800                 |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF A800                 |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| FAS9500                  |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF A900                 |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF C80                  |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| FAS90                    |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF A90                  |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |
| AFF A1K                  |          |                       |         | Note 1 |                     |         |                      |                     |         |                  |         |

### Supported ASA MetroCluster IP expansion combinations

The following table shows the supported platform combinations for expanding an ASA system in a MetroCluster IP configuration:

| ASA                      |          | Eight-node DR group 2 |          |          |          |          |          |          |          |
|--------------------------|----------|-----------------------|----------|----------|----------|----------|----------|----------|----------|
|                          |          | ASA A150              | ASA C250 | ASA A250 | ASA C400 | ASA A400 | ASA C800 | ASA A800 | ASA A900 |
| Eight-node<br>DR group 1 | ASA A150 |                       |          |          |          |          |          |          |          |
|                          | ASA C250 |                       |          |          |          |          |          |          |          |
|                          | ASA A250 |                       |          |          |          |          |          |          |          |
|                          | ASA C400 |                       |          |          |          |          |          |          |          |
|                          | ASA A400 |                       |          |          |          |          |          |          |          |
|                          | ASA C800 |                       |          |          |          |          |          |          |          |
|                          | ASA A800 |                       |          |          |          |          |          |          |          |
|                          | ASA A900 |                       |          |          |          |          |          |          |          |

### Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

#### About this task

This task must be performed on each MetroCluster site.

#### Steps

1. To prevent automatic support case generation, send an Autosupport message to indicate the upgrade is underway.
  - a. Issue the following command:

```
system node autosupport invoke -node * -type all -message "MAINT=10h
Upgrading <old-model> to <new-model>
```

This example specifies a 10 hour maintenance window. You might want to allow additional time, depending on your plan.

If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- b. Repeat the command on the partner cluster.

## Considerations for VLANs when adding a new DR group

- The following VLAN considerations apply when expanding a MetroCluster IP configuration:

Certain platforms use a VLAN for the MetroCluster IP interface. By default, each of the two ports use a different VLAN: 10 and 20.

If supported, you can also specify a different (non-default) VLAN higher than 100 (between 101 and 4095) using the `-vlan-id` parameter in the `metrocluster configuration-settings interface create` command.

The following platforms do **not** support the `-vlan-id` parameter:

- FAS8200 and AFF A300
- AFF A320
- FAS9000 and AFF A700
- AFF C800, ASA C800, AFF A800 and ASA A800

All other platforms support the `-vlan-id` parameter.

The default and valid VLAN assignments depend on whether the platform supports the `-vlan-id` parameter:

### Platforms that support `-vlan-id`

Default VLAN:

- When the `-vlan-id` parameter is not specified, the interfaces are created with VLAN 10 for the "A" ports and VLAN 20 for the "B" ports.
- The VLAN specified must match the VLAN selected in the RCF.

Valid VLAN ranges:

- Default VLAN 10 and 20
- VLANs 101 and higher (between 101 and 4095)

### Platforms that do not support `-vlan-id`

Default VLAN:

- Not applicable. The interface does not require a VLAN to be specified on the MetroCluster interface. The switch port defines the VLAN that is used.

Valid VLAN ranges:

- All VLANs not explicitly excluded when generating the RCF. The RCF alerts you if the VLAN is invalid.

- Both DR groups use the same VLANs when you expand from a four-node to an eight-node MetroCluster configuration.
- If both DR groups cannot be configured using the same VLAN, you must upgrade the DR group that doesn't support the `vlan-id` parameter to use a VLAN that is supported by the other DR group.

## Verifying the health of the MetroCluster configuration

You must verify the health and connectivity of the MetroCluster configuration prior to performing the expansion.

### Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

a. Check whether the system is multipathed:

```
node run -node <node-name> sysconfig -a
```

b. Check for any health alerts on both clusters:

```
system health alert show
```

c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

d. Perform a MetroCluster check:

```
metrocluster check run
```

e. Display the results of the MetroCluster check:

```
metrocluster check show
```

f. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

g. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

2. Verify that the cluster is healthy:

```
cluster show
```

```
cluster_A::> cluster show
Node Health Eligibility

node_A_1 true true
node_A_2 true true

cluster_A::>
```

3. Verify that all cluster ports are up:

```
network port show -ipSpace Cluster
```

```
cluster_A::> network port show -ipspace Cluster
```

```
Node: node_A_1-old
```

| Port | IPspace | Broadcast | Domain | Link | MTU  | Speed (Mbps)<br>Admin/Oper | Health<br>Status |
|------|---------|-----------|--------|------|------|----------------------------|------------------|
| e0a  | Cluster | Cluster   |        | up   | 9000 | auto/10000                 | healthy          |
| e0b  | Cluster | Cluster   |        | up   | 9000 | auto/10000                 | healthy          |

```
Node: node_A_2-old
```

| Port | IPspace | Broadcast | Domain | Link | MTU  | Speed (Mbps)<br>Admin/Oper | Health<br>Status |
|------|---------|-----------|--------|------|------|----------------------------|------------------|
| e0a  | Cluster | Cluster   |        | up   | 9000 | auto/10000                 | healthy          |
| e0b  | Cluster | Cluster   |        | up   | 9000 | auto/10000                 | healthy          |

```
4 entries were displayed.
```

```
cluster_A::>
```

#### 4. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

```
cluster_A::> network interface show -vserver cluster
```

| Current Is | Logical            | Status     | Network           | Current  |       |
|------------|--------------------|------------|-------------------|----------|-------|
| Vserver    | Interface          | Admin/Oper | Address/Mask      | Node     | Port  |
| Home       |                    |            |                   |          |       |
|            | -----              | -----      | -----             | -----    | ----- |
| Cluster    |                    |            |                   |          |       |
| true       | node_A_1-old_clus1 | up/up      | 169.254.209.69/16 | node_A_1 | e0a   |
| true       | node_A_1-old_clus2 | up/up      | 169.254.49.125/16 | node_A_1 | e0b   |
| true       | node_A_2-old_clus1 | up/up      | 169.254.47.194/16 | node_A_2 | e0a   |
| true       | node_A_2-old_clus2 | up/up      | 169.254.19.183/16 | node_A_2 | e0b   |

```
4 entries were displayed.
```

```
cluster_A::>
```

5. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

```

cluster_A::> network interface show -vserver Cluster -fields auto-revert

 Logical
Vserver Interface Auto-revert
----- -
Cluster
 node_A_1-old_clus1
 true
 node_A_1-old_clus2
 true
 node_A_2-old_clus1
 true
 node_A_2-old_clus2
 true

 4 entries were displayed.

cluster_A::>

```

## Removing the configuration from monitoring applications

If the existing configuration is monitored with the MetroCluster Tiebreaker software, the ONTAP Mediator or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the monitoring software prior to upgrade.

### Steps

1. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

| If you are using...      | Use this procedure...                                                                                                    |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Tiebreaker               | <a href="#">Removing MetroCluster Configurations.</a>                                                                    |
| Mediator                 | Issue the following command from the ONTAP prompt:<br><br><pre>metrocluster configuration-settings mediator remove</pre> |
| Third-party applications | Refer to the product documentation.                                                                                      |

2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.  
Refer to the documentation for the application.

## Preparing the new controller modules

You must prepare the four new MetroCluster nodes and install the correct ONTAP version.

### About this task

This task must be performed on each of the new nodes:

- node\_A\_3-new
- node\_A\_4-new
- node\_B\_3-new
- node\_B\_4-new

In these steps, you clear the configuration on the nodes and clear the mailbox region on new drives.

### Steps

1. Rack the new controllers.
2. Cable the new MetroCluster IP nodes to the IP switches as shown in [Cable the IP switches](#).
3. Configure the MetroCluster IP nodes using the following procedures:
  - a. [Gather required information](#)
  - b. [Restore system defaults on a controller module](#)
  - c. [Verify the ha-config state of components](#)
  - d. [Manually assign drives for pool 0 \(ONTAP 9.4 and later\)](#)
4. From Maintenance mode, issue the halt command to exit Maintenance mode, and then issue the `boot_ontap` command to boot the system and get to cluster setup.

Do not complete the cluster wizard or node wizard at this time.

## Upgrade RCF files

If you are installing new switch firmware, you must install the switch firmware before upgrading the RCF file.

### About this task

This procedure disrupts traffic on the switch where the RCF file is upgraded. Traffic will resume once the new RCF file is applied.

### Steps

1. Verify the health of the configuration.
  - a. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::*> metrocluster check run
```

The operation runs in the background.

- a. After the `metrocluster check run` operation completes, run `metrocluster check show` to

view the results.

After approximately five minutes, the following results are displayed:

```

::*> metrocluster check show

Component Result

nodes ok
lifs ok
config-replication ok
aggregates ok
clusters ok
connections ok
volumes ok
7 entries were displayed.
```

b. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id 38
```

c. Verify that there are no health alerts:

```
system health alert show
```

2. Prepare the IP switches for the application of the new RCF files.

Follow the steps for your switch vendor:

- [Reset the Broadcom IP switch to factory defaults](#)
- [Reset the Cisco IP switch to factory defaults](#)
- [Reset the NVIDIA IP SN2100 switch to factory defaults](#)

3. Download and install the IP RCF file, depending on your switch vendor.



Update the switches in the following order: Switch\_A\_1, Switch\_B\_1, Switch\_A\_2, Switch\_B\_2

- [Download and install the Broadcom IP RCF files](#)
- [Download and install the Cisco IP RCF files](#)
- [Download and install the NVIDIA IP RCF files](#)



If you have an L2 shared or L3 network configuration, you might need to adjust the ISL ports on the intermediate/customer switches. The switch port mode might change from 'access' to 'trunk' mode. Only proceed to upgrade the second switch pair (A\_2, B\_2) if the network connectivity between switches A\_1 and B\_1 is fully operational and the network is healthy.

## Join the new nodes to the clusters

You must add the four new MetroCluster IP nodes to the existing MetroCluster configuration.

### About this task

You must perform this task on both clusters.

### Steps

1. Add the new MetroCluster IP nodes to the existing MetroCluster configuration.
  - a. Join the first new MetroCluster IP node (node\_A\_1-new) to the existing MetroCluster IP configuration.

```
Welcome to the cluster setup wizard.
```

```
You can enter the following commands at any time:
```

```
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and
"exit" or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
```

```
You can return to cluster setup at any time by typing "cluster
setup".
```

```
To accept a default or omit a question, do not enter a value.
```

```
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
```

```
Enabling AutoSupport can significantly speed problem determination
and
resolution, should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
```

```
Type yes to confirm and continue {yes}: yes
```

```
Enter the node management interface port [e0M]: 172.17.8.93
```

```
172.17.8.93 is not a valid port.
```

```
The physical port that is connected to the node management network.
Examples of
node management ports are "e4a" or "e0M".
```

```
You can type "back", "exit", or "help" at any question.
```

```
Enter the node management interface port [e0M]:
Enter the node management interface IP address: 172.17.8.93
Enter the node management interface netmask: 255.255.254.0
Enter the node management interface default gateway: 172.17.8.1
A node management interface on port e0M with IP address 172.17.8.93
has been created.
```

Use your web browser to complete cluster setup by accessing  
<https://172.17.8.93>

Otherwise, press Enter to complete cluster setup using the command  
line  
interface:

```
Do you want to create a new cluster or join an existing cluster?
{create, join}:
join
```

Existing cluster interface configuration found:

| Port | MTU  | IP              | Netmask     |
|------|------|-----------------|-------------|
| e0c  | 9000 | 169.254.148.217 | 255.255.0.0 |
| e0d  | 9000 | 169.254.144.238 | 255.255.0.0 |

```
Do you want to use this configuration? {yes, no} [yes]: yes
```

```
.
.
.
```

b. Join the second new MetroCluster IP node (node\_A\_2-new) to the existing MetroCluster IP configuration.

2. Repeat these steps to join node\_B\_1-new and node\_B\_2-new to cluster\_B.

## Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates

You must create cluster peering LIFs, create the MetroCluster interfaces on the new MetroCluster IP nodes.

### About this task

- The home port used in the examples are platform-specific. You should use the home port specific to your MetroCluster IP node platform.
- Review the information in [Considerations for VLANs when adding a new DR group](#) before performing this task.

## Steps

1. On the new MetroCluster IP nodes, configure the intercluster LIFs using the following procedures:

[Configuring intercluster LIFs on dedicated ports](#)

[Configuring intercluster LIFs on shared data ports](#)

2. On each site, verify that cluster peering is configured:

```
cluster peer show
```

The following example shows the cluster peering configuration on cluster\_A:

```
cluster_A:> cluster peer show
Peer Cluster Name Cluster Serial Number Availability
Authentication

cluster_B 1-80-000011 Available ok
```

The following example shows the cluster peering configuration on cluster\_B:

```
cluster_B:> cluster peer show
Peer Cluster Name Cluster Serial Number Availability
Authentication

cluster_A 1-80-000011 Available ok
cluster_B::>
```

3. Create the DR group for the MetroCluster IP nodes:

```
metrocluster configuration-settings dr-group create -partner-cluster
```

For more information on the MetroCluster configuration settings and connections, see the following:

[Considerations for MetroCluster IP configurations](#)

[Creating the DR group](#)

```
cluster_A::> metrocluster configuration-settings dr-group create
-partner-cluster
cluster_B -local-node node_A_1-new -remote-node node_B_1-new
[Job 259] Job succeeded: DR Group Create is successful.
cluster_A::>
```

#### 4. Verify that the DR group was created.

```
metrocluster configuration-settings dr-group show
```

```
cluster_A::> metrocluster configuration-settings dr-group show

DR Group ID Cluster Node DR Partner
Node

1 cluster_A
node_A_1-old node_B_1-old
node_A_2-old node_B_2-old
cluster_B
node_B_1-old node_A_1-old
node_B_2-old node_A_2-old
2 cluster_A
node_A_1-new node_B_1-new
node_A_2-new node_B_2-new
cluster_B
node_B_1-new node_A_1-new
node_B_2-new node_A_2-new

8 entries were displayed.

cluster_A::>
```

#### 5. Configure the MetroCluster IP interfaces for the newly joined MetroCluster IP nodes:



- Do not use 169.254.17.x or 169.254.18.x IP addresses when you create MetroCluster IP interfaces to avoid conflicts with system auto-generated interface IP addresses in the same range.
- If supported, you can specify a different (non-default) VLAN higher than 100 (between 101 and 4095) using the `-vlan-id` parameter in the `metrocluster configuration-settings interface create` command. Refer to [Considerations for VLANs when adding a new DR group](#) for supported platform information.
- You can configure the MetroCluster IP interfaces from either cluster.

```
metrocluster configuration-settings interface create -cluster-name
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_1-new -home-port ela -address
172.17.26.10 -netmask 255.255.255.0
[Job 260] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_1-new -home-port elb -address
172.17.27.10 -netmask 255.255.255.0
[Job 261] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_2-new -home-port ela -address
172.17.26.11 -netmask 255.255.255.0
[Job 262] Job succeeded: Interface Create is successful.
```

```
cluster_A::> :metrocluster configuration-settings interface create
-cluster-name cluster_A -home-node node_A_2-new -home-port elb -address
172.17.27.11 -netmask 255.255.255.0
[Job 263] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_1-new -home-port ela -address
172.17.26.12 -netmask 255.255.255.0
[Job 264] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_1-new -home-port elb -address
172.17.27.12 -netmask 255.255.255.0
[Job 265] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_2-new -home-port ela -address
172.17.26.13 -netmask 255.255.255.0
[Job 266] Job succeeded: Interface Create is successful.
```

```
cluster_A::> metrocluster configuration-settings interface create
-cluster-name cluster_B -home-node node_B_2-new -home-port elb -address
172.17.27.13 -netmask 255.255.255.0
[Job 267] Job succeeded: Interface Create is successful.
```

## 6. Verify the MetroCluster IP interfaces are created:

```
metrocluster configuration-settings interface show
```

```
cluster_A::>metrocluster configuration-settings interface show
```

```

DR
Config
Group Cluster Node Network Address Netmask Gateway
State

1 cluster_A
 node_A_1-old
 Home Port: e1a
 172.17.26.10 255.255.255.0 -
completed
 Home Port: e1b
 172.17.27.10 255.255.255.0 -
completed
 node_A_2-old
 Home Port: e1a
 172.17.26.11 255.255.255.0 -
completed
 Home Port: e1b
 172.17.27.11 255.255.255.0 -
completed
 cluster_B
 node_B_1-old
 Home Port: e1a
 172.17.26.13 255.255.255.0 -
completed
 Home Port: e1b
 172.17.27.13 255.255.255.0 -
completed
 node_B_1-old
 Home Port: e1a
 172.17.26.12 255.255.255.0 -
completed
 Home Port: e1b
 172.17.27.12 255.255.255.0 -
completed
2 cluster_A
 node_A_3-new
 Home Port: e1a
 172.17.28.10 255.255.255.0 -
completed
 Home Port: e1b
 172.17.29.10 255.255.255.0 -
completed
 node_A_3-new

```



```

1 cluster_A
 node_A_1-old
 Home Port: e1a
 172.17.28.10 172.17.28.11 HA Partner
completed
 Home Port: e1a
 172.17.28.10 172.17.28.12 DR Partner
completed
 Home Port: e1a
 172.17.28.10 172.17.28.13 DR Auxiliary
completed
 Home Port: e1b
 172.17.29.10 172.17.29.11 HA Partner
completed
 Home Port: e1b
 172.17.29.10 172.17.29.12 DR Partner
completed
 Home Port: e1b
 172.17.29.10 172.17.29.13 DR Auxiliary
completed
 node_A_2-old
 Home Port: e1a
 172.17.28.11 172.17.28.10 HA Partner
completed
 Home Port: e1a
 172.17.28.11 172.17.28.13 DR Partner
completed
 Home Port: e1a
 172.17.28.11 172.17.28.12 DR Auxiliary
completed
 Home Port: e1b
 172.17.29.11 172.17.29.10 HA Partner
completed
 Home Port: e1b
 172.17.29.11 172.17.29.13 DR Partner
completed
 Home Port: e1b
 172.17.29.11 172.17.29.12 DR Auxiliary
completed

DR Source Destination
Group Cluster Node Network Address Network Address Partner Type
Config State

1 cluster_B

```

```

node_B_2-old
 Home Port: ela
 172.17.28.13 172.17.28.12 HA Partner
completed
 Home Port: ela
 172.17.28.13 172.17.28.11 DR Partner
completed
 Home Port: ela
 172.17.28.13 172.17.28.10 DR Auxiliary
completed
 Home Port: elb
 172.17.29.13 172.17.29.12 HA Partner
completed
 Home Port: elb
 172.17.29.13 172.17.29.11 DR Partner
completed
 Home Port: elb
 172.17.29.13 172.17.29.10 DR Auxiliary
completed
node_B_1-old
 Home Port: ela
 172.17.28.12 172.17.28.13 HA Partner
completed
 Home Port: ela
 172.17.28.12 172.17.28.10 DR Partner
completed
 Home Port: ela
 172.17.28.12 172.17.28.11 DR Auxiliary
completed
 Home Port: elb
 172.17.29.12 172.17.29.13 HA Partner
completed
 Home Port: elb
 172.17.29.12 172.17.29.10 DR Partner
completed
 Home Port: elb
 172.17.29.12 172.17.29.11 DR Auxiliary
completed

DR Source Destination
Group Cluster Node Network Address Network Address Partner Type
Config State

2 cluster_A
 node_A_1-new**

```

```

Home Port: e1a
172.17.26.10 172.17.26.11 HA Partner
completed

Home Port: e1a
172.17.26.10 172.17.26.12 DR Partner
completed

Home Port: e1a
172.17.26.10 172.17.26.13 DR Auxiliary
completed

Home Port: e1b
172.17.27.10 172.17.27.11 HA Partner
completed

Home Port: e1b
172.17.27.10 172.17.27.12 DR Partner
completed

Home Port: e1b
172.17.27.10 172.17.27.13 DR Auxiliary
completed

node_A_2-new
Home Port: e1a
172.17.26.11 172.17.26.10 HA Partner
completed

Home Port: e1a
172.17.26.11 172.17.26.13 DR Partner
completed

Home Port: e1a
172.17.26.11 172.17.26.12 DR Auxiliary
completed

Home Port: e1b
172.17.27.11 172.17.27.10 HA Partner
completed

Home Port: e1b
172.17.27.11 172.17.27.13 DR Partner
completed

Home Port: e1b
172.17.27.11 172.17.27.12 DR Auxiliary
completed

DR
Group Cluster Node Source Destination Partner Type
Config State

2 cluster_B
 node_B_2-new
 Home Port: e1a

```

```

172.17.26.13 172.17.26.12 HA Partner
completed
Home Port: e1a
172.17.26.13 172.17.26.11 DR Partner
completed
Home Port: e1a
172.17.26.13 172.17.26.10 DR Auxiliary
completed
Home Port: e1b
172.17.27.13 172.17.27.12 HA Partner
completed
Home Port: e1b
172.17.27.13 172.17.27.11 DR Partner
completed
Home Port: e1b
172.17.27.13 172.17.27.10 DR Auxiliary
completed
node_B_1-new
Home Port: e1a
172.17.26.12 172.17.26.13 HA Partner
completed
Home Port: e1a
172.17.26.12 172.17.26.10 DR Partner
completed
Home Port: e1a
172.17.26.12 172.17.26.11 DR Auxiliary
completed
Home Port: e1b
172.17.27.12 172.17.27.13 HA Partner
completed
Home Port: e1b
172.17.27.12 172.17.27.10 DR Partner
completed
Home Port: e1b
172.17.27.12 172.17.27.11 DR Auxiliary
completed
48 entries were displayed.

cluster_A::>

```

#### 9. Verify disk auto-assignment and partitioning:

```
disk show -pool Pool1
```

```

cluster_A::> disk show -pool Pool1
 Usable Disk Container Container
Disk Size Shelf Bay Type Type Name
Owner

1.10.4 - 10 4 SAS remote -
node_B_2
1.10.13 - 10 13 SAS remote -
node_B_2
1.10.14 - 10 14 SAS remote -
node_B_1
1.10.15 - 10 15 SAS remote -
node_B_1
1.10.16 - 10 16 SAS remote -
node_B_1
1.10.18 - 10 18 SAS remote -
node_B_2
...
2.20.0 546.9GB 20 0 SAS aggregate aggr0_rha1_a1
node_a_1
2.20.3 546.9GB 20 3 SAS aggregate aggr0_rha1_a2
node_a_2
2.20.5 546.9GB 20 5 SAS aggregate rha1_a1_aggr1
node_a_1
2.20.6 546.9GB 20 6 SAS aggregate rha1_a1_aggr1
node_a_1
2.20.7 546.9GB 20 7 SAS aggregate rha1_a2_aggr1
node_a_2
2.20.10 546.9GB 20 10 SAS aggregate rha1_a1_aggr1
node_a_1
...
43 entries were displayed.

cluster_A::>

```

## 10. Mirror the root aggregates:

```
storage aggregate mirror -aggregate aggr0_node_A_1-new
```



You must complete this step on each MetroCluster IP node.

```

cluster_A::> aggr mirror -aggregate aggr0_node_A_1-new

Info: Disks would be added to aggregate "aggr0_node_A_1-new"on node
"node_A_1-new"
 in the following manner:

 Second Plex

 RAID Group rg0, 3 disks (block checksum, raid_dp)

Physical Usable
Size Position Disk Type Size

- dparity 4.20.0 SAS -
- parity 4.20.3 SAS -
- data 4.20.1 SAS 546.9GB
558.9GB

 Aggregate capacity available forvolume use would be 467.6GB.

Do you want to continue? {y|n}: y

cluster_A::>

```

11. Verify that the root aggregates are mirrored:

```
storage aggregate show
```

```

cluster_A::> aggr show

Aggregate Size Available Used% State #Vols Nodes RAID
Status

aggr0_node_A_1-old
 349.0GB 16.84GB 95% online 1 node_A_1-old
raid_dp,
mirrored,
normal

```

```

aggr0_node_A_2-old
 349.0GB 16.84GB 95% online 1 node_A_2-old
raid_dp,

mirrored,

normal
aggr0_node_A_1-new
 467.6GB 22.63GB 95% online 1 node_A_1-new
raid_dp,

mirrored,

normal
aggr0_node_A_2-new
 467.6GB 22.62GB 95% online 1 node_A_2-new
raid_dp,

mirrored,

normal
aggr_data_a1
 1.02TB 1.01TB 1% online 1 node_A_1-old
raid_dp,

mirrored,

normal
aggr_data_a2
 1.02TB 1.01TB 1% online 1 node_A_2-old
raid_dp,

mirrored,

```

## Finalizing the addition of the new nodes

You must incorporate the new DR group into the MetroCluster configuration and create mirrored data aggregates on the new nodes.

### Steps

1. Refresh the MetroCluster configuration:

- a. Enter advanced privilege mode:

```
set -privilege advanced
```

- b. Refresh the MetroCluster configuration on any of the newly added nodes:

| If your MetroCluster configuration has...             | Then do this...                                                                                                                                                                                                |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Multiple data aggregates                              | From any node's prompt, run:<br><br><code>metrocluster configure &lt;node-name&gt;</code>                                                                                                                      |
| A single mirrored data aggregate at one or both sites | From any node's prompt, configure the MetroCluster with the <code>-allow-with-one-aggregate true</code> parameter:<br><br><code>metrocluster configure -allow-with-one-aggregate true &lt;node-name&gt;</code> |

c. Reboot each of the new nodes:

```
node reboot -node <node_name> -inhibit-takeover true
```



You don't need to reboot the nodes in a specific order, but you should wait until one node is fully booted and all connections are established before rebooting the next node.

d. Return to admin privilege mode:

```
set -privilege admin
```

2. Create mirrored data aggregates on each of the new MetroCluster nodes:

```
storage aggregate create -aggregate <aggregate-name> -node <node-name>
-diskcount <no-of-disks> -mirror true
```



- You must create at least one mirrored data aggregate per site. It is recommended to have two mirrored data aggregates per site on MetroCluster IP nodes to host the MDV volumes, however a single aggregate per site is supported (but not recommended). It is acceptable that one site of the MetroCluster has a single mirrored data aggregate and the other site has more than one mirrored data aggregate.
- Aggregate names must be unique across the MetroCluster sites. This means that you cannot have two different aggregates with the same name on site A and site B.

The following example shows the creation of an aggregate on `node_A_1-new`.

```
cluster_A::> storage aggregate create -aggregate data_a3 -node node_A_1-
new -diskcount 10 -mirror t
```

```
Info: The layout for aggregate "data_a3" on node "node_A_1-new" would
be:
```

```
First Plex
```

```
RAID Group rg0, 5 disks (block checksum, raid_dp)
```

```
Usable
```

```

Physical
Size Position Disk Type Size

- dparity 5.10.15 SAS -
- parity 5.10.16 SAS -
- data 5.10.17 SAS 546.9GB
547.1GB data 5.10.18 SAS 546.9GB
558.9GB data 5.10.19 SAS 546.9GB
558.9GB

```

Second Plex

RAID Group rg0, 5 disks (block checksum, raid\_dp)

```

Usable
Physical
Size Position Disk Type Size

- dparity 4.20.17 SAS -
- parity 4.20.14 SAS -
- data 4.20.18 SAS 546.9GB
547.1GB data 4.20.19 SAS 546.9GB
547.1GB data 4.20.16 SAS 546.9GB
547.1GB

```

Aggregate capacity available for volume use would be 1.37TB.

Do you want to continue? {y|n}: y

[Job 440] Job succeeded: DONE

cluster\_A::>

3. Verify that the nodes are added to their DR group.

```
cluster_A::*> metrocluster node show
```

| DR Group | Cluster   | Node         | Configuration State | DR Mirroring | Mode   |
|----------|-----------|--------------|---------------------|--------------|--------|
| 1        | cluster_A | node_A_1-old | configured          | enabled      | normal |
|          |           | node_A_2-old | configured          | enabled      | normal |
|          | cluster_B | node_B_1-old | configured          | enabled      | normal |
|          |           | node_B_2-old | configured          | enabled      | normal |
| 2        | cluster_A | node_A_3-new | configured          | enabled      | normal |
|          |           | node_A_4-new | configured          | enabled      | normal |
|          | cluster_B | node_B_3-new | configured          | enabled      | normal |
|          |           | node_B_4-new | configured          | enabled      | normal |

8 entries were displayed.

```
cluster_A::*>
```

#### 4. Move the MDV\_CRS volumes in advanced privilege mode.

##### a. Display the volumes to identify the MDV volumes:

If you have a single mirrored data aggregate per site then move both the MDV volumes to this single aggregate. If you have two or more mirrored data aggregates, then move each MDV volume to a different aggregate.

If you are expanding a four-node MetroCluster configuration to a permanent eight-node configuration, you should move one of the MDV volumes to the new DR group.

The following example shows the MDV volumes in the `volume show` output:

```

cluster_A::> volume show
Vserver Volume Aggregate State Type Size
Available Used%

...

cluster_A MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_A
 aggr_b1 - RW -
- -
cluster_A MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_B
 aggr_b2 - RW -
- -
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_A
 aggr_a1 online RW 10GB
9.50GB 0%
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
 aggr_a2 online RW 10GB
9.50GB 0%
...
11 entries were displayed.mple

```

b. Set the advanced privilege level:

```
set -privilege advanced
```

c. Move the MDV volumes, one at a time:

```
volume move start -volume <mdv-volume> -destination-aggregate <aggr-on-new-
node> -vserver <svm-name>
```

The following example shows the command and output for moving "MDV\_CRS\_d6b0b313ff5611e9837100a098544e51\_A" to aggregate "data\_a3" on "node\_A\_3".

```

cluster_A::*> vol move start -volume
MDV_CRS_d6b0b313ff5611e9837100a098544e51_A -destination-aggregate
data_a3 -vserver cluster_A

Warning: You are about to modify the system volume
 "MDV_CRS_d6b0b313ff5611e9837100a098544e51_A". This might
cause severe
 performance or stability problems. Do not proceed unless
directed to
 do so by support. Do you want to proceed? {y|n}: y
[Job 494] Job is queued: Move
"MDV_CRS_d6b0b313ff5611e9837100a098544e51_A" in Vserver "cluster_A"
to aggregate "data_a3". Use the "volume move show -vserver cluster_A
-volume MDV_CRS_d6b0b313ff5611e9837100a098544e51_A" command to view
the status of this operation.

```

d. Use the volume show command to check that the MDV volume has been successfully moved:

```
volume show <mdv-name>
```

The following output shows that the MDV volume has been successfully moved.

```

cluster_A::*> vol show MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
Vserver Volume Aggregate State Type Size
Available Used%

cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
 aggr_a2 online RW 10GB
9.50GB 0%

```

5. Move epsilon from an old node to a new node:

a. Identify which node currently has epsilon:

```
cluster show -fields epsilon
```

```
cluster_B::*> cluster show -fields epsilon
node epsilon

node_A_1-old true
node_A_2-old false
node_A_3-new false
node_A_4-new false
4 entries were displayed.
```

- b. Set epsilon to false on the old node (node\_A\_1-old):

```
cluster modify -node <old-node> -epsilon false*
```

- c. Set epsilon to true on the new node (node\_A\_3-new):

```
cluster modify -node <new-node> -epsilon true
```

- d. Verify that epsilon has moved to the correct node:

```
cluster show -fields epsilon
```

```
cluster_A::*> cluster show -fields epsilon
node epsilon

node_A_1-old false
node_A_2-old false
node_A_3-new true
node_A_4-new false
4 entries were displayed.
```

6. If your system supports end-to-end encryption, you can [Enable end-to-end encryption](#) on the new DR group.

## Remove a DR group from a MetroCluster configuration

Beginning with ONTAP 9.8, you can remove a disaster recovery (DR) group from an eight-node MetroCluster configuration to create a four-node MetroCluster configuration.



You use these steps during the transition and system refresh workflows.

### Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.

- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

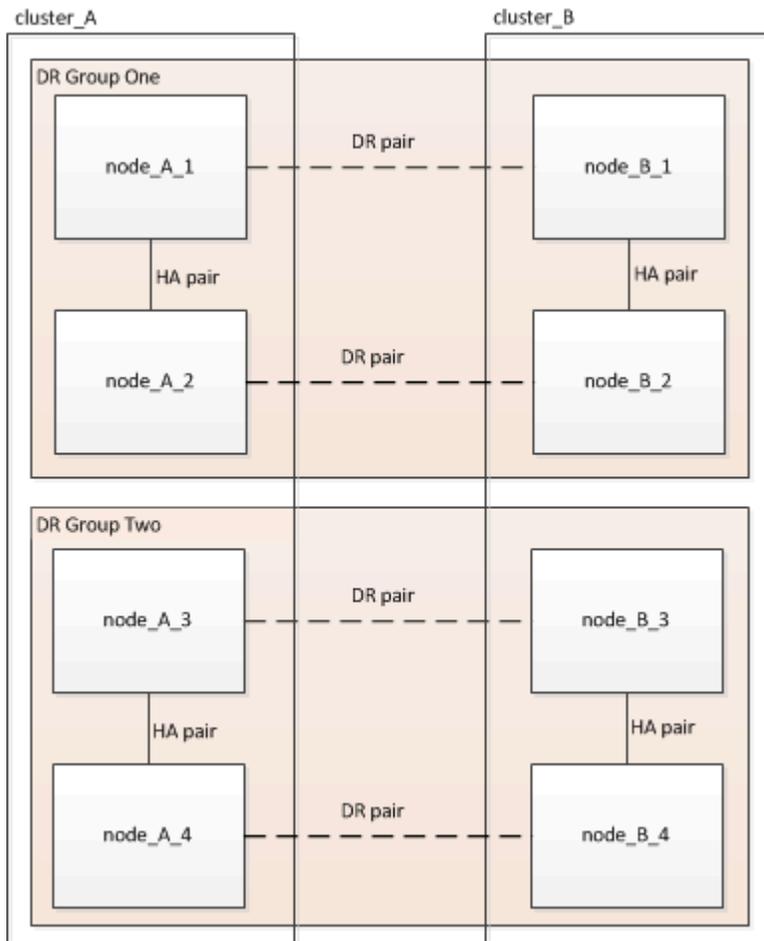
- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Remove the DR group nodes from each cluster

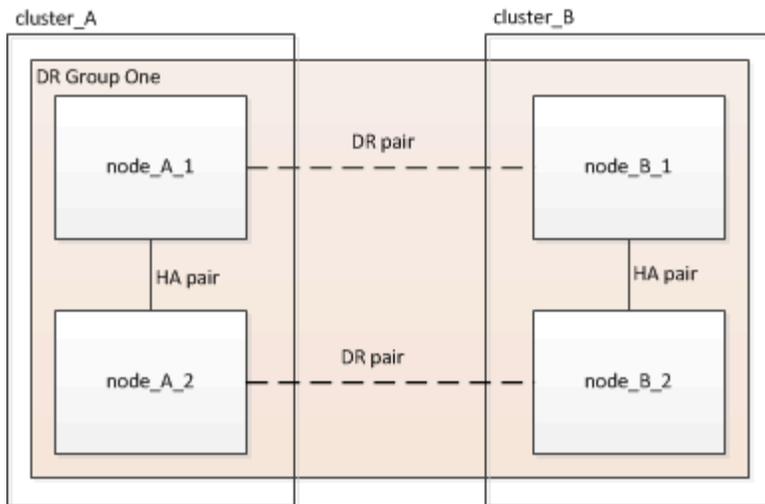
This procedure is supported beginning with ONTAP 9.8. For systems running ONTAP 9.7 or earlier, see the Knowledge Base article: [How to remove a DR group from a MetroCluster configuration](#).

### About this task

An eight-node configuration includes eight-nodes organized as two four-node DR groups.



When you remove a DR group, four nodes remain in the configuration.



### Before you begin

- You must perform this step on both clusters.
- The `metrocluster remove-dr-group` command is supported only on ONTAP 9.8 and later.

### Steps

1. Prepare to remove the DR group if you have not already.
  - a. Move all data volumes to another DR group.
  - b. If the DR group to be removed has load-sharing mirror volumes, re-create all load-sharing mirror volumes in another DR group and delete them from the DR group to be removed.
  - c. Move all MDV\_CRS metadata volumes to another DR group by following the [Moving a metadata volume in MetroCluster configurations](#) procedure.
  - d. Delete all MDV\_aud metadata volumes that might exist in the DR group to be removed.
  - e. Delete all data aggregates in the DR group to be removed:

```
ClusterA::> storage aggregate show -node ClusterA-01, ClusterA-02
-fields aggregate ,node
ClusterA::> aggr delete -aggregate aggregate_name
ClusterB::> storage aggregate show -node ClusterB-01, ClusterB-02
-fields aggregate ,node
ClusterB::> aggr delete -aggregate aggregate_name
```



Root aggregates are not deleted.

- f. Migrate all NAS data LIFs that you use for NFS and CIFS (SMB) to home nodes in another DR group.

```
network interface show -home-node <old_node>
```

```
network interface migrate -vserver <svm_name> -lif <data_lif> -destination
-node <new_node> -destination-port <port>
```

9. Move the data LIFs to the new home node in another DR group.

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-node
```

```
<new_node> -home-port <port>
```

h. Migrate the cluster management LIF to a home node in another DR group.

```
network interface show -role cluster-mgmt
```

```
network interface modify -vserver <svm-name> -lif <cluster_mgmt> -home-node
<new_node> -home-port <port-id>
```



- Node management and inter-cluster LIFs are not migrated. Create new node management and inter-cluster LIFs on nodes of the DR group as required.
- You cannot migrate or move FCP interfaces used for block access (SAN) between the nodes. Create new FCP interfaces as needed.
- iSCSI SAN LIFs need to be down before the home node and home port can be updated.

i. Transfer epsilon to a node in another DR group if required.

```
ClusterA::> set advanced
ClusterA::*> cluster show
Move epsilon if needed
ClusterA::*> cluster modify -node nodename -epsilon false
ClusterA::*> cluster modify -node nodename -epsilon true

ClusterB::> set advanced
ClusterB::*> cluster show
ClusterB::*> cluster modify -node nodename -epsilon false
ClusterB::*> cluster modify -node nodename -epsilon true
ClusterB::*> set admin
```

2. Identify and remove the DR group.

a. Identify the correct DR group for removal:

```
metrocluster node show
```

b. Remove the DR group nodes:

```
metrocluster remove-dr-group -dr-group-id 1
```

The following example shows the removal of the DR group configuration on cluster\_A.

## Example

```
cluster_A::~*>

Warning: Nodes in the DR group that are removed from the
MetroCluster
 configuration will lose their disaster recovery
protection.

 Local nodes "node_A_1-FC, node_A_2-FC"will be removed
from the
 MetroCluster configuration. You must repeat the
operation on the
 partner cluster "cluster_B"to remove the remote nodes in
the DR group.
Do you want to continue? {y|n}: y

Info: The following preparation steps must be completed on the
local and partner
 clusters before removing a DR group.

 1. Move all data volumes to another DR group.
 2. Move all MDV_CRS metadata volumes to another DR group.
 3. Delete all MDV_aud metadata volumes that may exist in
the DR group to
 be removed.
 4. Delete all data aggregates in the DR group to be
removed. Root
 aggregates are not deleted.
 5. Migrate all data LIFs to home nodes in another DR group.
 6. Migrate the cluster management LIF to a home node in
another DR group.
 Node management and inter-cluster LIFs are not migrated.
 7. Transfer epsilon to a node in another DR group.

 The command is vetoed if the preparation steps are not
completed on the
 local and partner clusters.
Do you want to continue? {y|n}: y
[Job 513] Job succeeded: Remove DR Group is successful.

cluster_A::~*>
```

3. Repeat the previous step on the partner cluster.

4. Disable storage failover on the nodes of the old DR group:

```
storage failover modify -node <node-name> -enable false
```

5. If you are in a MetroCluster IP configuration, perform the following steps to delete the remote plexes of the root aggregates and remove disk ownership on the nodes of the old DR group.

These steps need to be performed for both nodes in the HA pair at each site.

a. Display the remote plexes of root aggregates on the nodes in the DR group that is to be deleted:

```
storage aggregate plex show -aggregate <root_aggr_name> -pool 1
```

b. Delete the remote plexes:

```
storage aggregate plex delete -aggregate <root_aggr_name> -plex
<plex_from_previous_step>
```

c. Identify the remote disks owned by the nodes in the DR group.

The commands you use depend on whether you are using partitioned/shared disks or whole disks:



Use a comma-separated list in the `-owner <node_names>` field to specify the node names in the DR group that is to be deleted.

**Partitioned/shared disks:**

i. Set the privilege level to advanced:

```
set advanced
```

ii. Display the remote disks:

```
storage disk show -pool Pool1 -owner <node_names> -partition
-ownership
```

**Whole disks:**

i. Set the privilege level to advanced:

```
set advanced
```

ii. Display the remote disks:

```
storage disk show -pool Pool1 -owner <node_names>
```

d. Disable disk auto assignment:

```
disk option modify -node <node_names_in_the_DR_group_to_be_deleted>
-autoassign off
```

e. Remove ownership of pool1 disks on each DR group node to be deleted. Perform these steps on each

node to be removed.

- i. Go to the nodeshell:

```
run -node <node_name>
```

- ii. Identify the pool1 disks:

```
aggr status -s
```

All spare disks are displayed, including the pool0 and pool1 spare disks owned by the node.

- iii. Remove disk ownership for each pool1 spare disk:

```
disk remove_ownership <disk_name>
```

For partitioned disks, remove partition ownership and then remove the container disk ownership.

6. If you are in a MetroCluster IP configuration, remove the MetroCluster connections on the nodes of the old DR group.

These commands can be issued from either cluster and apply to the entire DR group spanning both of the clusters.

- a. Disconnect the connections:

```
metrocluster configuration-settings connection disconnect -dr-group-id
<dr_group_id>
```

### Example

```
cluster_A::*> metrocluster configuration-settings connection
disconnect -dr-group-id 1
```

```
Warning: For the nodes in the DR group 1, this command will
remove the existing connections that are used to mirror NV logs
and access remote storage.
```

```
Do you want to continue? {y|n}: y
```

```
Warning: Before proceeding with disconnect, you must verify the
following:
```

```
1. Unmirrored aggregates do not have disks in remote
plexes.
```

```
2. Aggregates are not mirrored.
```

```
3. No disks are assigned in Pool1.
```

```
4. Storage failover is not enabled.
```

```
Follow the "MetroCluster Installation and Configuration
guide" for detailed instructions to verify this.
```

```
Do you want to continue? {y|n}: y
```

b. Delete the MetroCluster interfaces on the nodes of the old DR group:



This step must be repeated on each node of the DR group.

```
metrocluster configuration-settings interface delete
```

c. Delete the old DR group's configuration.

```
metrocluster configuration-settings dr-group delete
```

7. Unjoin the nodes in the old DR group.

Perform this step on each cluster.

a. Set the advanced privilege level:

```
set -privilege advanced
```

b. Unjoin the node:

```
cluster unjoin -node <node-name>
```

Repeat this step for the other local node in the old DR group.

c. Set the admin privilege level:

```
set -privilege admin
```

8. Check that cluster HA is enabled in the new DR group. If required, re-enable cluster HA:

```
cluster ha modify -configured true
```

Perform this step on each cluster.

9. Halt, power down, and remove the old controller modules and storage shelves.

## Where to find additional information

You can learn more about MetroCluster configuration and operation.

### MetroCluster and miscellaneous information

| Information                                | Subject                                                                      |
|--------------------------------------------|------------------------------------------------------------------------------|
| <a href="#">MetroCluster Documentation</a> | <ul style="list-style-type: none"><li>All MetroCluster information</li></ul> |

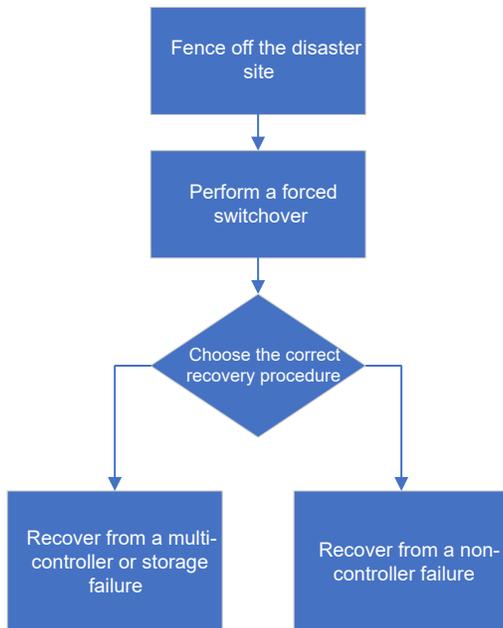
|                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Fabric-attached MetroCluster installation and configuration</p>     | <ul style="list-style-type: none"> <li>• Fabric-attached MetroCluster architecture</li> <li>• Cabling the configuration</li> <li>• Configuring the FC-to-SAS bridges</li> <li>• Configuring the FC switches</li> <li>• Configuring the MetroCluster in ONTAP</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <p>Stretch MetroCluster installation and configuration</p>             | <ul style="list-style-type: none"> <li>• Stretch MetroCluster architecture</li> <li>• Cabling the configuration</li> <li>• Configuring the FC-to-SAS bridges</li> <li>• Configuring the MetroCluster in ONTAP</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <p>MetroCluster management and disaster recovery</p>                   | <ul style="list-style-type: none"> <li>• Understanding the MetroCluster configuration</li> <li>• Switchover, healing and switchback</li> <li>• Disaster recovery</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <p>Maintain MetroCluster Components</p>                                | <ul style="list-style-type: none"> <li>• Guidelines for maintenance in a MetroCluster FC configuration</li> <li>• Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches</li> <li>• Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration</li> <li>• Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration</li> <li>• Replacing hardware at a disaster site in a fabric-attached or stretch MetroCluster FC configuration</li> <li>• Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration.</li> <li>• Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.</li> </ul> |
| <p>MetroCluster Upgrade, Transition, and Expansion</p>                 | <ul style="list-style-type: none"> <li>• Upgrading or refreshing a MetroCluster configuration</li> <li>• Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration</li> <li>• Expanding a MetroCluster configuration by adding additional nodes</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <p>MetroCluster Tiebreaker Software installation and configuration</p> | <ul style="list-style-type: none"> <li>• Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

|                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| <p><a href="#">ONTAP Hardware Systems Documentation</a></p> <div style="display: flex; align-items: center;">  <p>The standard storage shelf maintenance procedures can be used with MetroCluster IP configurations.</p> </div> | <ul style="list-style-type: none"> <li>• Hot-adding a disk shelf</li> <li>• Hot-removing a disk shelf</li> </ul>                |
| <p><a href="#">Copy-based transition</a></p>                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• Transitioning data from 7-Mode storage systems to clustered storage systems</li> </ul> |
| <p><a href="#">ONTAP concepts</a></p>                                                                                                                                                                                                                                                                                                        | <ul style="list-style-type: none"> <li>• How mirrored aggregates work</li> </ul>                                                |

# Recover from a disaster

## Workflow for disaster recovery

Use the workflow to perform disaster recovery.



## Performing a forced switchover after a disaster

If a disaster has occurred, there are steps you must perform on both the disaster cluster and the surviving cluster after the switchover to ensure safe and continued data service.

Determining if a disaster has occurred is done by:

- An administrator
- The MetroCluster Tiebreaker software, if it is configured
- The ONTAP Mediator software, if it is configured

### Fencing off the disaster site

After the disaster, if the disaster site nodes must be replaced, you must halt them to prevent the site from resuming service. Otherwise, you risk the possibility of data corruption if clients start accessing the nodes before the replacement procedure is completed.

#### Step

1. Halt the nodes at the disaster site and keep them powered down or at the LOADER prompt until directed to boot ONTAP:

```
system node halt -node disaster-site-node-name
```

If the disaster site nodes have been destroyed or cannot be halted, turn off power to the nodes and do not boot the replacement nodes until directed to in the recovery procedure.

## Performing a forced switchover

The switchover process, in addition to providing nondisruptive operations during testing and maintenance, enables you to recover from a site failure with a single command.

### Before you begin

- At least one of the surviving site nodes must be up and running before you perform the switchover.
- All previous configuration changes must be complete before performing a switchback operation.

This is to avoid competition with the negotiated switchover or switchback operation.



SnapMirror and SnapVault configurations are deleted automatically.

### About this task

The `metrocluster switchover` command switches over the nodes in all DR groups in the MetroCluster configuration. For example, in an eight-node MetroCluster configuration, it switches over the nodes in both DR groups.

### Steps

1. Perform the switchover by running the following command at the surviving site:

```
metrocluster switchover -forced-on-disaster true
```



The operation can take a period of minutes to complete. You can verify progress using the `metrocluster operation show` command.

2. Answer `y` when prompted to continue with the switchover.
3. Verify that the switchover was completed successfully by running the `metrocluster operation show` command.

```
mccl1A::> metrocluster operation show
 Operation: switchover
 Start time: 10/4/2012 19:04:13
 State: in-progress
 End time: -
 Errors:

mccl1A::> metrocluster operation show
 Operation: switchover
 Start time: 10/4/2012 19:04:13
 State: successful
 End time: 10/4/2012 19:04:22
 Errors: -
```

If the switchover is vetoed, you have the option of reissuing the `metrocluster switchover-forced-on-disaster true` command with the `--override-vetoes` option. If you use this optional parameter, the system overrides any soft vetoes that prevented the switchover.

### After you finish

SnapMirror relationships need to be reestablished after switchover.

## Output for the storage aggregate plex show command is indeterminate after a MetroCluster switchover

When you run the `storage aggregate plex show` command after a MetroCluster switchover, the status of plex0 of the switched over root aggregate is indeterminate and is displayed as failed. During this time, the switched over root is not updated. The actual status of this plex can only be determined after the MetroCluster healing phase.

## Accessing volumes in NVFAIL state after a switchover

After a switchover, you must clear the NVFAIL state by resetting the `-in-nvfailed-state` parameter of the `volume modify` command to remove the restriction of clients to access data.

### Before you begin

The database or file system must not be running or trying to access the affected volume.

### About this task

Setting the `-in-nvfailed-state` parameter requires advanced-level privilege.

### Step

1. Recover the volume by using the `volume modify` command with the `-in-nvfailed-state` parameter set to false.

### After you finish

For instructions about examining database file validity, see the documentation for your specific database software.

If your database uses LUNs, review the steps to make the LUNs accessible to the host after an NVRAM failure.

### Related information

[Monitoring and protecting database validity by using NVFAIL](#)

## Choosing the correct recovery procedure

After a failure in a MetroCluster configuration, you must select the correct recovery procedure. Use the following table and examples to select the appropriate recovery procedure.

This information in this table assumes that the installation or transition is complete, meaning that the `metrocluster configure` command ran successfully.

| Scope of failures at disaster site | Procedure |
|------------------------------------|-----------|
|------------------------------------|-----------|

|                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• No hardware failure (for example, a power failure)</li> </ul>                                                                                                                                                   | <a href="#">Recovering from a non-controller failure</a>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <ul style="list-style-type: none"> <li>• No controller module failure</li> <li>• Other hardware has failed</li> </ul>                                                                                                                                    | <a href="#">Recovering from a non-controller failure</a>                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <ul style="list-style-type: none"> <li>• Single controller module failure or failure of FRU components within the controller module</li> <li>• Drives have not failed</li> </ul>                                                                         | <p>If a failure is limited to a single controller module, you must use the controller module FRU replacement procedure for the platform model. In a four or eight-node MetroCluster configuration, such a failure is isolated to the local HA pair.</p> <p><b>Note:</b> The controller module FRU replacement procedure can be used in a two-node MetroCluster configuration if there are no drive or other hardware failures.</p> <p><a href="#">ONTAP Hardware Systems Documentation</a></p> |
| <ul style="list-style-type: none"> <li>• Single controller module failure or failure of FRU components within the controller module</li> <li>• Drives have failed</li> </ul>                                                                             | <a href="#">Recovering from a multi-controller or storage failure</a>                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <ul style="list-style-type: none"> <li>• Single controller module failure or failure of FRU components within the controller module</li> <li>• Drives have not failed</li> <li>• Additional hardware outside the controller module has failed</li> </ul> | <a href="#">Recovering from a multi-controller or storage failure</a><br><p>You should skip all steps for drive assignment.</p>                                                                                                                                                                                                                                                                                                                                                                |
| <ul style="list-style-type: none"> <li>• Multiple controller module failure (with or without additional failures) within a DR group</li> </ul>                                                                                                           | <a href="#">Recovering from a multi-controller or storage failure</a>                                                                                                                                                                                                                                                                                                                                                                                                                          |

## Controller module failure scenarios during MetroCluster installation

Responding to a controller module failure during the MetroCluster configuration procedure depends on whether the `metrocluster configure` command successfully completed.

- If the `metrocluster configure` command was not yet run, or failed, you must restart the MetroCluster software configuration procedure from the beginning with a replacement controller module.



You must be sure to perform the steps in [Restoring system defaults on a controller module](#) on each controller (including the replacement controller) to verify that the previous configuration is removed.

- If the `metrocluster configure` command successfully completed and then the controller module failed, use the previous table to determine the correct recovery procedure.

## Controller module failure scenarios during MetroCluster FC-to-IP transition

The recovery procedure can be used if a site failure occurs during transition. However, it can only be used if the configuration is a stable mixed configuration, with the FC DR group and IP DR group both fully configured. The output of the `metrocluster node show` command should show both DR groups with all eight nodes.



If the failure occurred during transition when the nodes are in the process of being added or removed, you must contact technical support.

## Controller module failure scenarios in eight-node MetroCluster configurations

Failure scenarios:

- [Single controller module failures in a single DR group](#)
- [Two controller module failures in a single DR group](#)
- [Single controller module failures in separate DR groups](#)
- [Three controller module failures spread across the DR groups](#)

### Single controller module failures in a single DR group

In this case the failure is limited to an HA pair.

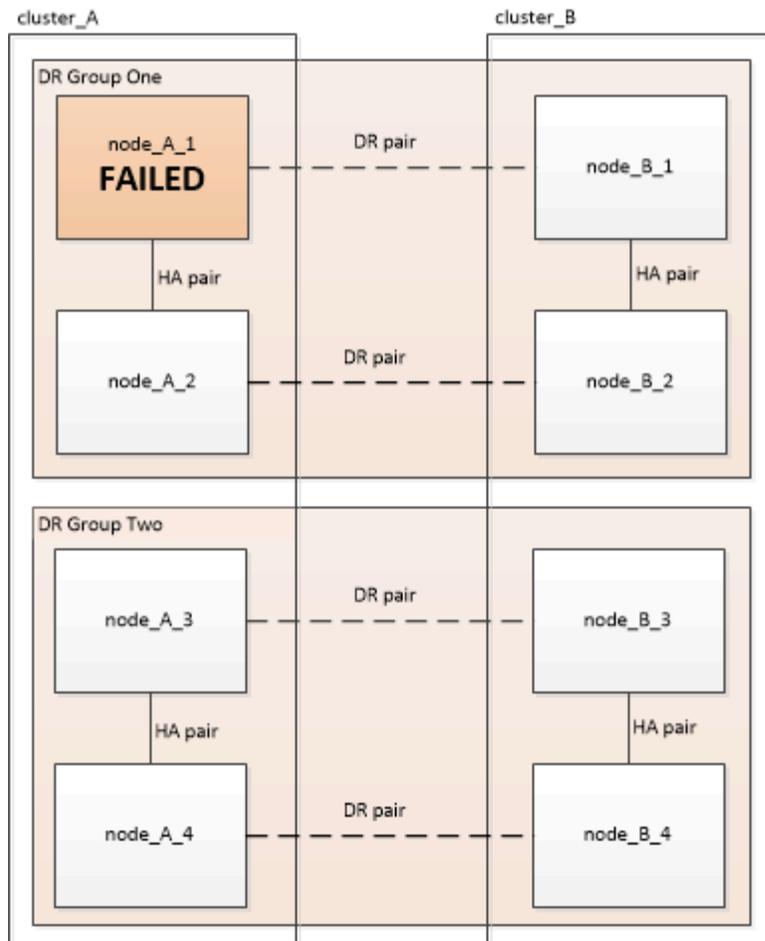
- If no storage requires replacement, you can use the controller module FRU replacement procedure for the platform model.

[ONTAP Hardware Systems Documentation](#)

- If storage requires replacement, you can use the multi-controller module recovery procedure.

[Recovering from a multi-controller or storage failure](#)

This scenario applies to four-node MetroCluster configurations also.

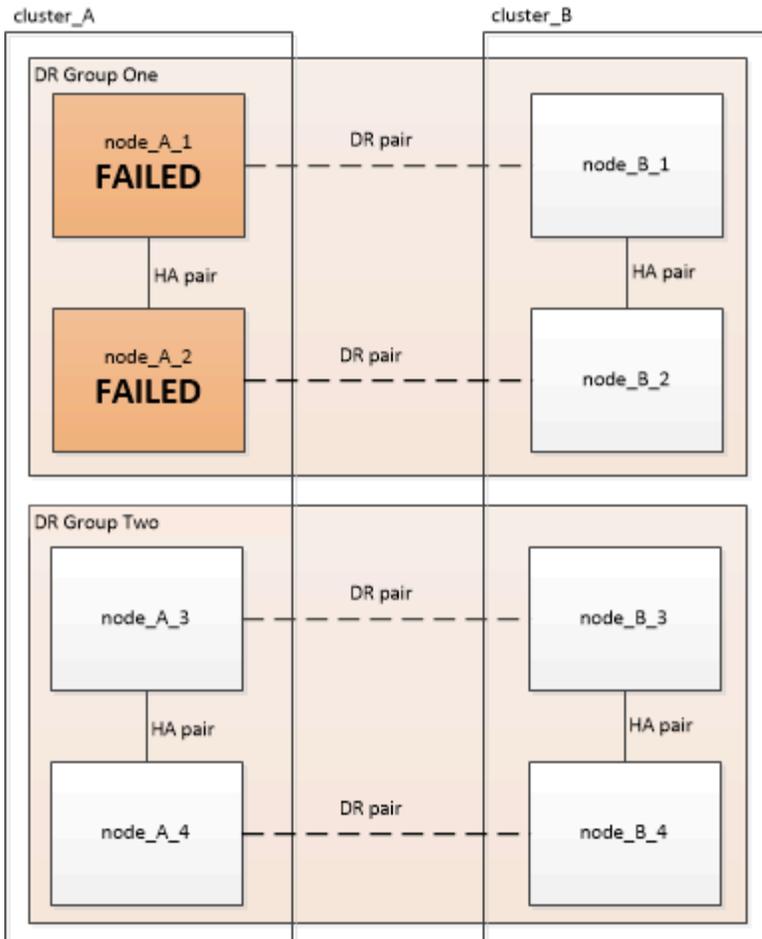


### Two controller module failures in a single DR group

In this case the failure requires a switchover. You can use the multi-controller module failure recovery procedure.

#### [Recovering from a multi-controller or storage failure](#)

This scenario applies to four-node MetroCluster configurations also.



### Single controller module failures in separate DR groups

In this case the failure is limited to separate HA pairs.

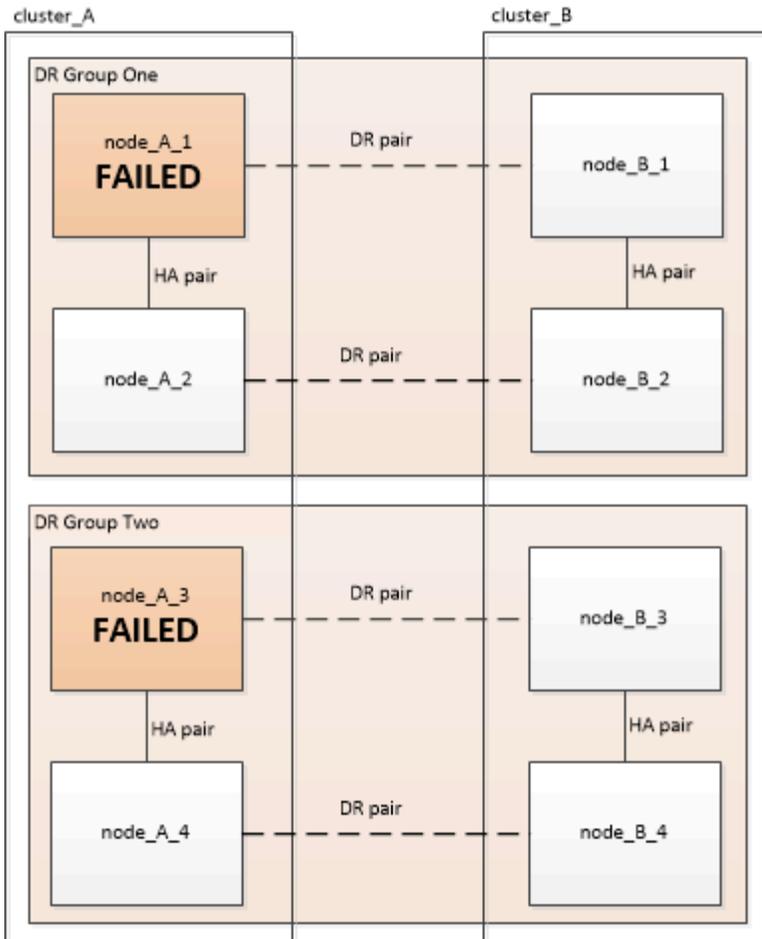
- If no storage requires replacement, you can use the controller module FRU replacement procedure for the platform model.

The FRU replacement procedure is performed twice, once for each failed controller module.

[ONTAP Hardware Systems Documentation](#)

- If storage requires replacement, you can use the multi-controller module recovery procedure.

[Recovering from a multi-controller or storage failure](#)



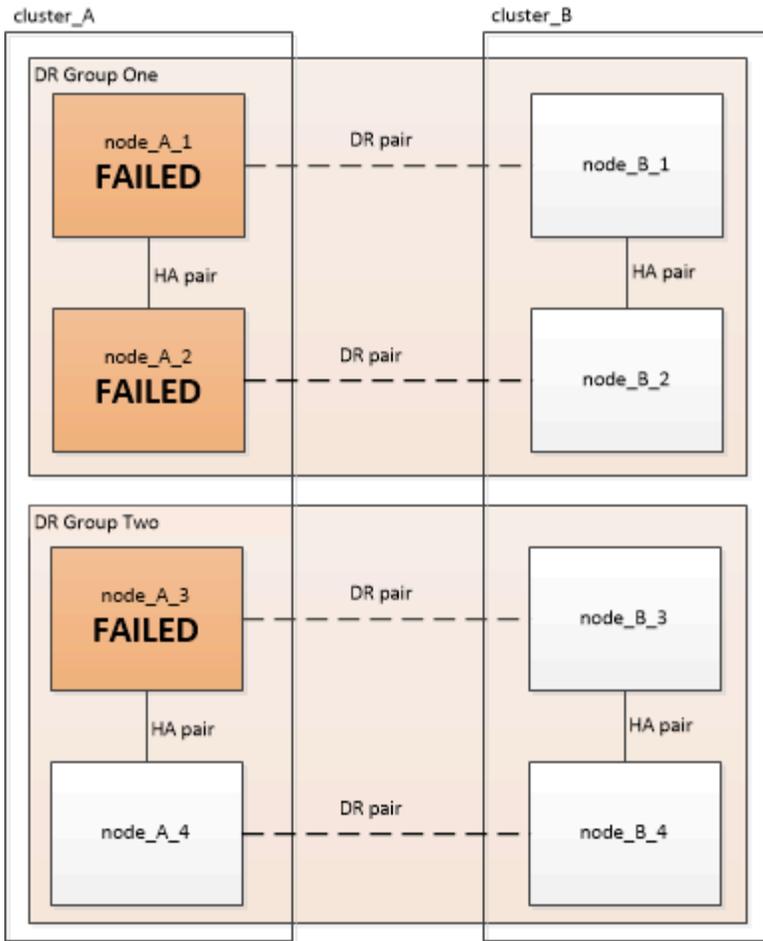
### Three controller module failures spread across the DR groups

In this case the failure requires a switchover. You can use the multi-controller module failure recovery procedure for DR Group One.

#### [Recovering from a multi-controller or storage failure](#)

You can use the platform-specific controller module FRU replacement procedure for DR Group Two.

[ONTAP Hardware Systems Documentation](#)



## Controller module failure scenarios in two-node MetroCluster configurations

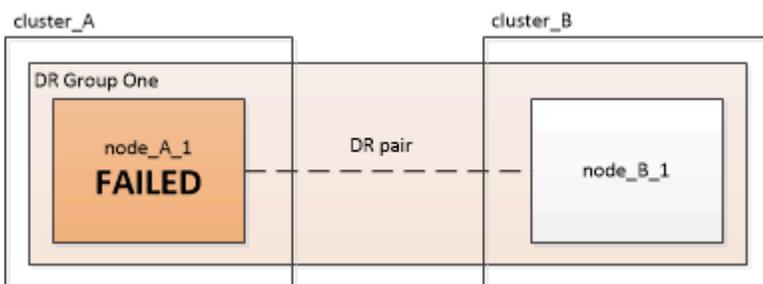
The procedure you use depends on the extent of the failure.

- If no storage requires replacement, you can use the controller module FRU replacement procedure for the platform model.

[ONTAP Hardware Systems Documentation](#)

- If storage requires replacement, you can use the multi-controller module recovery procedure.

[Recovering from a multi-controller or storage failure](#)



# Recover from a multi-controller or storage failure

## Recovering from a multi-controller or storage failure

If the controller failure extends to all controller modules on one side of a DR group in a MetroCluster configuration (including a single controller in a two-node MetroCluster configuration), or storage has been replaced, you must replace the equipment and reassign ownership of drives to recover from the disaster.

Verify that you have checked and performed the following tasks before using this procedure:

- Review the available recovery procedures before deciding to use this procedure.

[Choosing the correct recovery procedure](#)

- Confirm that console logging is enabled on your devices.

[Enable console logging](#)

- Ensure that the disaster site is fenced off.

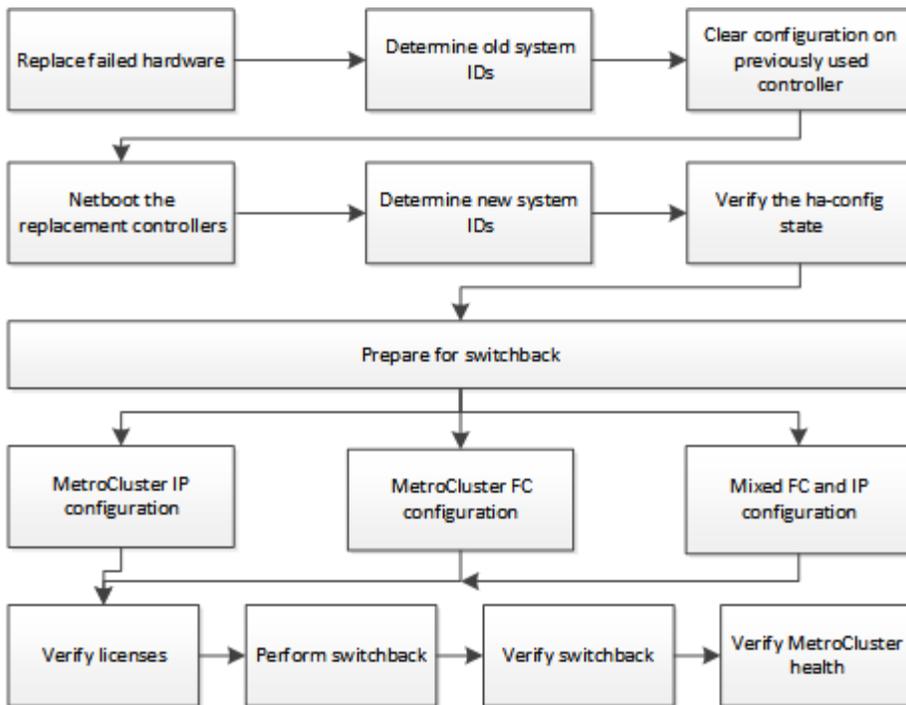
[Fencing off the disaster site.](#)

- Verify that switchover was performed.

[Performing a forced switchover.](#)

- Verify that the replacement drives and the controller modules are new and must not have been assigned ownership previously.
- The examples in this procedure show two or four-node configurations. If you have an eight-node configuration (two DR groups), you must take into account any failures and perform the required recovery task on the additional controller modules.

This procedure uses the following workflow:



This procedure can be used when performing recovery on a system that was in mid-transition when the failure occurred. In that case, you must perform the appropriate steps when preparing for switchback, as indicated in the procedure.

## Enable console logging

Enable console logging on your devices before proceeding to replace hardware and boot new controllers.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Replace hardware and boot new controllers

If hardware components have to be replaced, you must replace them using their individual hardware replacement and installation guides.

### Replace hardware at the disaster site

#### Before you begin

The storage controllers must be powered off or remain halted (showing the LOADER prompt).

## Steps

1. Replace the components as necessary.



In this step, you replace and cable the components exactly as they were cabled prior to the disaster. You must not power up the components.

| If you are replacing...                        | Perform these steps...                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Using these guides...                                                                                                   |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| FC switches in a MetroCluster FC configuration | <ol style="list-style-type: none"> <li>a. Install the new switches.</li> <li>b. Cable the ISL links.<br/>Do not power on the FC switches at this time.</li> </ol>                                                                                                                                                                                                                                                                                                                             | <a href="#">Maintain MetroCluster Components</a>                                                                        |
| IP switches in a MetroCluster IP configuration | <ol style="list-style-type: none"> <li>a. Install the new switches.</li> <li>b. Cable the ISL links.<br/>Do not power on the IP switches at this time.</li> </ol>                                                                                                                                                                                                                                                                                                                             | <a href="#">MetroCluster IP installation and configuration: Differences among the ONTAP MetroCluster configurations</a> |
| Disk shelves                                   | <ol style="list-style-type: none"> <li>a. Install the disk shelves and disks.               <ul style="list-style-type: none"> <li>◦ Disk shelf stacks should be the same configuration as at the surviving site.</li> <li>◦ Disks can be the same size or larger, but must be of the same type (SAS or SATA).</li> </ul> </li> <li>b. Cable the disk shelves to adjacent shelves within the stack and to the FC-to-SAS bridge.<br/>Do not power on the disk shelves at this time.</li> </ol> | <a href="#">ONTAP Hardware Systems Documentation</a>                                                                    |
| SAS cables                                     | <ol style="list-style-type: none"> <li>a. Install the new cables.<br/>Do not power on the disk shelves at this time.</li> </ol>                                                                                                                                                                                                                                                                                                                                                               | <a href="#">ONTAP Hardware Systems Documentation</a>                                                                    |

|                                                             |                                                                                                                                                                                                                                                               |                                                                                                                                                               |
|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>FC-to-SAS bridges in a MetroCluster FC configuration</p> | <p>a. Install the FC-to-SAS bridges.<br/>b. Cable the FC-to-SAS bridges.</p> <p>Cable them to the FC switches or to the controller modules, depending on your MetroCluster configuration type.</p> <p>Do not power on the FC-to-SAS bridges at this time.</p> | <p><a href="#">Fabric-attached MetroCluster installation and configuration</a></p> <p><a href="#">Stretch MetroCluster installation and configuration</a></p> |
|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                      |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|
| Controller modules | <p>a. Install the new controller modules:</p> <ul style="list-style-type: none"> <li>◦ The controller modules must be the same model as those being replaced.</li> </ul> <p>For example, 8080 controller modules must be replaced with 8080 controller modules.</p> <ul style="list-style-type: none"> <li>◦ The controller modules must not have previously been part of either cluster within the MetroCluster configuration or any previously existing cluster configuration.</li> </ul> <p>If they were, you must set defaults and perform a “wipeconfig” process.</p> <ul style="list-style-type: none"> <li>◦ Ensure that all network interface cards (such as Ethernet or FC) are in the same slots used on the old controller modules.</li> </ul> <p>b. Cable the new controller modules exactly the same as the old ones.</p> <p>The ports connecting the controller module to the storage (either by connections to the IP or FC switches, FC-to-SAS bridges, or directly) should be the same as those used prior to the disaster.</p> <p>Do not power on the controller modules at this time.</p> | <a href="#">ONTAP Hardware Systems Documentation</a> |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------|

2. Verify that all components are cabled correctly for your configuration.

- [MetroCluster IP configuration](#)
- [MetroCluster fabric-attached configuration](#)

### Determine the system IDs and VLAN IDs of the old controller modules

After you have replaced all hardware at the disaster site, you must determine the system IDs of the replaced controller modules. You need the old system IDs when you reassign disks to the new controller modules. If the

systems are AFF A220, AFF A250, AFF A400, AFF A800, FAS2750, FAS500f, FAS8300, or FAS8700 models, you must also determine the VLAN IDs used by the MetroCluster IP interfaces.

**Before you begin**

All equipment at the disaster site must be powered off.

**About this task**

This discussion provides examples for two and four-node configurations. For eight-node configurations, you must account for any failures in the additional nodes on the second DR group.

For a two-node MetroCluster configuration, you can ignore references to the second controller module at each site.

The examples in this procedure are based on the following assumptions:

- Site A is the disaster site.
- node\_A\_1 has failed and is being completely replaced.
- node\_A\_2 has failed and is being completely replaced.

node\_A\_2 is present in a four-node MetroCluster configuration only.

- Site B is the surviving site.
- node\_B\_1 is healthy.
- node\_B\_2 is healthy.

node\_B\_2 is present in a four-node MetroCluster configuration only.

The controller modules have the following original system IDs:

| Number of nodes in MetroCluster configuration | Node     | Original system ID |
|-----------------------------------------------|----------|--------------------|
| Four                                          | node_A_1 | 4068741258         |
|                                               | node_A_2 | 4068741260         |
|                                               | node_B_1 | 4068741254         |
|                                               | node_B_2 | 4068741256         |
| Two                                           | node_A_1 | 4068741258         |
|                                               | node_B_1 | 4068741254         |

**Steps**

1. From the surviving site, display the system IDs of the nodes in the MetroCluster configuration.

| Number of nodes in MetroCluster configuration | Use this command |
|-----------------------------------------------|------------------|
|-----------------------------------------------|------------------|

|               |                                                                                                                         |
|---------------|-------------------------------------------------------------------------------------------------------------------------|
| Four or eight | <code>metrocluster node show -fields node-systemid,ha-partner-systemid,dr-partner-systemid,dr-auxiliary-systemid</code> |
| Two           | <code>metrocluster node show -fields node-systemid,dr-partner-systemid</code>                                           |

In this example for a four-node MetroCluster configuration, the following old system IDs are retrieved:

- Node\_A\_1: 4068741258
- Node\_A\_2: 4068741260

Disks owned by the old controller modules are still owned these system IDs.

```

metrocluster node show -fields node-systemid,ha-partner-systemid,dr-
partner-systemid,dr-auxiliary-systemid

dr-group-id cluster node node-systemid ha-partner-systemid
dr-partner-systemid dr-auxiliary-systemid

1 Cluster_A Node_A_1 4068741258 4068741260
4068741254 4068741256
1 Cluster_A Node_A_2 4068741260 4068741258
4068741256 4068741254
1 Cluster_B Node_B_1 - -
-
1 Cluster_B Node_B_2 - -
-
4 entries were displayed.

```

In this example for a two-node MetroCluster configuration, the following old system ID is retrieved:

- Node\_A\_1: 4068741258

Disks owned by the old controller module are still owned this system ID.

```

metrocluster node show -fields node-systemid,dr-partner-systemid

dr-group-id cluster node node-systemid dr-partner-systemid

1 Cluster_A Node_A_1 4068741258 4068741254
1 Cluster_B Node_B_1 - -
2 entries were displayed.

```

2. For MetroCluster IP configurations using ONTAP Mediator, get the IP address of ONTAP Mediator:

```
storage iscsi-initiator show -node * -label mediator
```

3. If the systems are AFF A220, AFF A400, FAS2750, FAS8300, or FAS8700 models, determine the VLAN IDs:

```
metrocluster interconnect show
```

The VLAN IDs are included in the adapter name shown in the Adapter column of the output.

In this example, the VLAN IDs are 120 and 130:

```
metrocluster interconnect show
```

| Node     | Partner  | Name | Type | Mirror Admin Status | Mirror Oper Status | Adapter | Type  | Status |
|----------|----------|------|------|---------------------|--------------------|---------|-------|--------|
| Node_A_1 | Node_A_2 | HA   |      | enabled             | online             | e0a-120 | iWARP | Up     |
|          |          |      |      |                     |                    | e0b-130 | iWARP | Up     |
|          | Node_B_1 | DR   |      | enabled             | online             | e0a-120 | iWARP | Up     |
|          |          |      |      |                     |                    | e0b-130 | iWARP | Up     |
|          | Node_B_2 | AUX  |      | enabled             | offline            | e0a-120 | iWARP | Up     |
|          |          |      |      |                     |                    | e0b-130 | iWARP | Up     |
| Node_A_2 | Node_A_1 | HA   |      | enabled             | online             | e0a-120 | iWARP | Up     |
|          |          |      |      |                     |                    | e0b-130 | iWARP | Up     |
|          | Node_B_2 | DR   |      | enabled             | online             | e0a-120 | iWARP | Up     |
|          |          |      |      |                     |                    | e0b-130 | iWARP | Up     |
|          | Node_B_1 | AUX  |      | enabled             | offline            | e0a-120 | iWARP | Up     |
|          |          |      |      |                     |                    | e0b-130 | iWARP | Up     |

12 entries were displayed.

### Isolate replacement drives from the surviving site (MetroCluster IP configurations)

You must isolate any replacement drives by taking down the MetroCluster iSCSI initiator connections from the surviving nodes.

#### About this task

This procedure is only required on MetroCluster IP configurations.

#### Steps

1. From either surviving node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).

2. Disconnect the iSCSI initiators on both surviving nodes in the DR group:

```
storage iscsi-initiator disconnect -node surviving-node -label *
```

This command must be issued twice, once for each of the surviving nodes.

The following example shows the commands for disconnecting the initiators on site B:

```
site_B::*> storage iscsi-initiator disconnect -node node_B_1 -label *
site_B::*> storage iscsi-initiator disconnect -node node_B_2 -label *
```

3. Return to the admin privilege level:

```
set -privilege admin
```

## Clear the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the existing configuration.

### Steps

1. If necessary, halt the node to display the `LOADER` prompt:

```
halt
```

2. At the `LOADER` prompt, set the environmental variables to default values:

```
set-defaults
```

3. Save the environment:

```
saveenv
```

4. At the `LOADER` prompt, launch the boot menu:

```
boot_ontap menu
```

5. At the boot menu prompt, clear the configuration:

```
wipeconfig
```

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

6. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond *yes* to the confirmation prompt.

### Netboot the new controller modules

If the new controller modules have a different version of ONTAP from the version on the surviving controller modules, you must netboot the new controller modules.

#### Before you begin

- You must have access to an HTTP server.
- You must have access to the NetApp Support Site to download the necessary system files for your platform and version of ONTAP software that is running on it.

[NetApp Support](#)

#### Steps

1. Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software from the software download section of the NetApp Support Site and store the `ontap-version_image.tgz` file on a web-accessible directory.
3. Go to the web-accessible directory and verify that the files you need are available.

| If the platform model is... | Then...                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| FAS/AFF8000 series systems  | <p>Extract the contents of the <code>ontap-version_image.tgz</code> file to the target directory: <code>tar -zxvf ontap-version_image.tgz</code></p> <p>NOTE: If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image.</p> <p>Your directory listing should contain a netboot folder with a kernel file: <code>netboot/kernel</code></p> |
| All other systems           | <p>Your directory listing should contain a netboot folder with a kernel file: <code>ontap-version_image.tgz</code></p> <p>You do not need to extract the <code>ontap-version_image.tgz</code> file.</p>                                                                                                                                                                          |

4. At the LOADER prompt, configure the netboot connection for a management LIF:

- If IP addressing is DHCP, configure the automatic connection:

```
ifconfig e0M -auto
```

- If IP addressing is static, configure the manual connection:

```
ifconfig e0M -addr=ip_addr -mask=netmask -gw=gateway
```

5. Perform the netboot.

- If the platform is an 80xx series system, use this command:

```
netboot http://web_server_ip/path_to_web-accessible_directory/netboot/kernel
```

- If the platform is any other system, use the following command:

```
netboot http://web_server_ip/path_to_web-accessible_directory/ontap-
version_image.tgz
```

6. From the boot menu, select option **(7) Install new software first** to download and install the new software image to the boot device.

Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". It applies to nondisruptive upgrades of software, not to upgrades of controllers.

7. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the URL of the image file: `http://web_server_ip/path_to_web-accessible_directory/ontap-version_image.tgz`

Enter username/password if applicable, or press Enter to continue.

8. Enter `n` to skip the backup recovery when you see a prompt similar to the following:

Do you want to restore the backup configuration now? {y|n} n

9. Reboot by entering `y` when you see a prompt similar to the following:

The node must be rebooted to start using the newly installed software.  
Do you want to reboot now? {y|n} y



You must reboot the node in order to use the newly installed software.

10. From the Boot menu, select **option 5** to enter Maintenance mode.
11. If you have a four-node MetroCluster configuration, repeat this procedure on the other new controller module.

### Determine the system IDs of the replacement controller modules

After you have replaced all hardware at the disaster site, you must determine the system ID of the newly installed storage controller module or modules.

#### About this task

You must perform this procedure with the replacement controller modules in Maintenance mode.

This section provides examples for two and four-node configurations. For two-node configurations, you can ignore references to the second node at each site. For eight-node configurations, you must account for the additional nodes on the second DR group. The examples make the following assumptions:

- Site A is the disaster site.
- node\_A\_1 has been replaced.
- node\_A\_2 has been replaced.

Present only in four-node MetroCluster configurations.

- Site B is the surviving site.
- node\_B\_1 is healthy.
- node\_B\_2 is healthy.

Present only in four-node MetroCluster configurations.

The examples in this procedure use controllers with the following system IDs:

| Number of nodes in MetroCluster configuration | Node     | Original system ID | New system ID | Will pair with this node as DR partner |
|-----------------------------------------------|----------|--------------------|---------------|----------------------------------------|
| Four                                          | node_A_1 | 4068741258         | 1574774970    | node_B_1                               |
|                                               | node_A_2 | 4068741260         | 1574774991    | node_B_2                               |
|                                               | node_B_1 | 4068741254         | unchanged     | node_A_1                               |
|                                               | node_B_2 | 4068741256         | unchanged     | node_A_2                               |
| Two                                           | node_A_1 | 4068741258         | 1574774970    | node_B_1                               |
|                                               | node_B_1 | 4068741254         | unchanged     | node_A_1                               |



In a four-node MetroCluster configuration, the system determines DR partnerships by pairing the node with the lowest system ID at site\_A and the node with the lowest system ID at site\_B. Because the system IDs change, the DR pairs might be different after the controller replacements are completed than they were prior to the disaster.

In the preceding example:

- node\_A\_1 (1574774970) will be paired with node\_B\_1 (4068741254)
- node\_A\_2 (1574774991) will be paired with node\_B\_2 (4068741256)

### Steps

1. With the node in Maintenance mode, display the local system ID of the node from each node: `disk show`

In the following example, the new local system ID is 1574774970:

```
*> disk show
Local System ID: 1574774970
...
```

2. On the second node, repeat the previous step.



This step is not required in a two-node MetroCluster configuration.

In the following example, the new local system ID is 1574774991:

```
*> disk show
Local System ID: 1574774991
...
```

### Verify the ha-config state of components

In a MetroCluster configuration, the ha-config state of the controller module and chassis components must be set to "mcc" or "mcc-2n" so they boot up properly.

#### Before you begin

The system must be in Maintenance mode.

#### About this task

This task must be performed on each new controller module.

#### Steps

1. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The correct HA state depends on your MetroCluster configuration.

| Number of controllers in the MetroCluster configuration | HA state for all components should be... |
|---------------------------------------------------------|------------------------------------------|
| Eight- or four-node MetroCluster FC configuration       | mcc                                      |
| Two-node MetroCluster FC configuration                  | mcc-2n                                   |
| MetroCluster IP configuration                           | mccip                                    |

2. If the displayed system state of the controller is not correct, set the HA state for the controller module:

| Number of controllers in the MetroCluster configuration | Command |
|---------------------------------------------------------|---------|
|---------------------------------------------------------|---------|

|                                                   |                                    |
|---------------------------------------------------|------------------------------------|
| Eight- or four-node MetroCluster FC configuration | ha-config modify controller mcc    |
| Two-node MetroCluster FC configuration            | ha-config modify controller mcc-2n |
| MetroCluster IP configuration                     | ha-config modify controller mccip  |

- If the displayed system state of the chassis is not correct, set the HA state for the chassis:

| Number of controllers in the MetroCluster configuration | Command                         |
|---------------------------------------------------------|---------------------------------|
| Eight- or four-node MetroCluster FC configuration       | ha-config modify chassis mcc    |
| Two-node MetroCluster FC configuration                  | ha-config modify chassis mcc-2n |
| MetroCluster IP configuration                           | ha-config modify chassis mccip  |

- Repeat these steps on the other replacement node.

### Determine if end-to-end encryption was enabled on the original systems

You should verify if the original systems were configured for end-to-end encryption.

#### Step

- Run the following command from the surviving site:

```
metrocluster node show -fields is-encryption-enabled
```

If encryption is enabled, the following output is displayed:

```
1 cluster_A node_A_1 true
1 cluster_A node_A_2 true
1 cluster_B node_B_1 true
1 cluster_B node_B_2 true
4 entries were displayed.
```



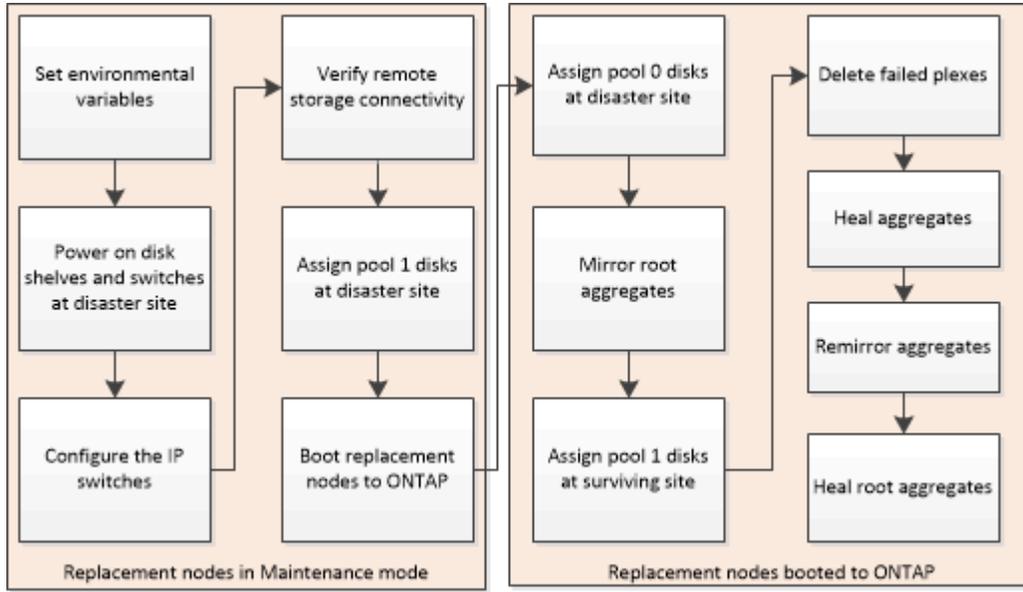
Refer to [Configure end-to-end encryption](#) for supported systems.

### Prepare for switchback in a MetroCluster IP configuration

#### Prepare for switchback in a MetroCluster IP configuration

You must perform certain tasks in order to prepare the MetroCluster IP configuration for the switchback operation.

## About this task



## Setting required environmental variables in MetroCluster IP configurations

In MetroCluster IP configurations, you must retrieve the IP address of the MetroCluster interfaces on the Ethernet ports, and then use them to configure the interfaces on the replacement controller modules.

### About this task

- This task is required only in MetroCluster IP configurations.
- Commands in this task are performed from the cluster prompt of the surviving site and from the LOADER prompt of the nodes at the disaster site.
- Certain platforms use a VLAN for the MetroCluster IP interface. By default, each of the two ports use a different VLAN: 10 and 20.

If supported, you can also specify a different (non-default) VLAN higher than 100 (between 101 and 4095) using the `vlan-id` parameter.

The following platforms do **not** support the `vlan-id` parameter:

- FAS8200 and AFF A300
- AFF A320
- FAS9000 and AFF A700
- AFF C800, ASA C800, AFF A800 and ASA A800

All other platforms support the `vlan-id` parameter.

- The nodes in these examples have the following IP addresses for their MetroCluster IP connections:



These examples are for an AFF A700 or FAS9000 system. The interfaces vary by platform model.

| Node     | Port | IP address   |
|----------|------|--------------|
| node_A_1 | e5a  | 172.17.26.10 |
|          | e5b  | 172.17.27.10 |
| node_A_2 | e5a  | 172.17.26.11 |
|          | e5b  | 172.17.27.11 |
| node_B_1 | e5a  | 172.17.26.13 |
|          | e5b  | 172.17.27.13 |
| node_B_2 | e5a  | 172.17.26.12 |
|          | e5b  | 172.17.27.12 |

The following table summarizes the relationships between the nodes and each node's MetroCluster IP addresses.

| Node                                                   | HA partner                                             | DR partner                                             | DR auxiliary partner                                   |
|--------------------------------------------------------|--------------------------------------------------------|--------------------------------------------------------|--------------------------------------------------------|
| node_A_1<br>• e5a: 172.17.26.10<br>• e5b: 172.17.27.10 | node_A_2<br>• e5a: 172.17.26.11<br>• e5b: 172.17.27.11 | node_B_1<br>• e5a: 172.17.26.13<br>• e5b: 172.17.27.13 | node_B_2<br>• e5a: 172.17.26.12<br>• e5b: 172.17.27.12 |
| node_A_2<br>• e5a: 172.17.26.11<br>• e5b: 172.17.27.11 | node_A_1<br>• e5a: 172.17.26.10<br>• e5b: 172.17.27.10 | node_B_2<br>• e5a: 172.17.26.12<br>• e5b: 172.17.27.12 | node_B_1<br>• e5a: 172.17.26.13<br>• e5b: 172.17.27.13 |
| node_B_1<br>• e5a: 172.17.26.13<br>• e5b: 172.17.27.13 | node_B_2<br>• e5a: 172.17.26.12<br>• e5b: 172.17.27.12 | node_A_1<br>• e5a: 172.17.26.10<br>• e5b: 172.17.27.10 | node_A_2<br>• e5a: 172.17.26.11<br>• e5b: 172.17.27.11 |
| node_B_2<br>• e5a: 172.17.26.12<br>• e5b: 172.17.27.12 | node_B_1<br>• e5a: 172.17.26.13<br>• e5b: 172.17.27.13 | node_A_2<br>• e5a: 172.17.26.11<br>• e5b: 172.17.27.11 | node_A_1<br>• e5a: 172.17.26.10<br>• e5b: 172.17.27.10 |

- The MetroCluster bootarg values you set depend on whether your new system uses shared cluster/HA ports or shared MetroCluster/HA ports. Use the following information to determine the ports for your system.

### Shared cluster/HA ports

The systems listed in the following table use shared cluster/HA ports:

| AFF and ASA systems                                                                                                                                                                                                 | FAS systems                                                                                 |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• AFF A20</li> <li>• AFF A30</li> <li>• AFF C30</li> <li>• AFF A50</li> <li>• AFF C60</li> <li>• AFF C80</li> <li>• AFF A70</li> <li>• AFF A90</li> <li>• AFF A1K</li> </ul> | <ul style="list-style-type: none"> <li>• FAS50</li> <li>• FAS70</li> <li>• FAS90</li> </ul> |

### Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

| AFF and ASA systems                                                                                                                                                                                                                                                                                                                                                      | FAS systems                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>• AFF A150, ASA A150</li> <li>• AFF A220</li> <li>• AFF C250, ASA C250</li> <li>• AFF A250, ASA A250</li> <li>• AFF A300</li> <li>• AFF A320</li> <li>• AFF C400, ASA C400</li> <li>• AFF A400, ASA A400</li> <li>• AFF A700</li> <li>• AFF C800, ASA C800</li> <li>• AFF A800, ASA A800</li> <li>• AFF A900, ASA A900</li> </ul> | <ul style="list-style-type: none"> <li>• FAS2750</li> <li>• FAS500f</li> <li>• FAS8200</li> <li>• FAS8300</li> <li>• FAS8700</li> <li>• FAS9000</li> <li>• FAS9500</li> </ul> |

### Steps

1. From the surviving site, gather the IP addresses of the MetroCluster interfaces on the disaster site:

```
metrocluster configuration-settings connection show
```

The required addresses are the DR Partner addresses shown in the **Destination Network Address** column.

The command output varies depending on whether your platform model uses shared cluster/HA ports or shared MetroCluster/HA ports.

### Systems using shared cluster/HA ports

```
cluster_B::*> metrocluster configuration-settings connection show
DR Source Destination
DR Source Destination
Group Cluster Node Network Address Network Address Partner Type
Config State

1 cluster_B
 node_B_1
 Home Port: e5a
 172.17.26.13 172.17.26.10 DR Partner
completed
 Home Port: e5a
 172.17.26.13 172.17.26.11 DR Auxiliary
completed
 Home Port: e5b
 172.17.27.13 172.17.27.10 DR Partner
completed
 Home Port: e5b
 172.17.27.13 172.17.27.11 DR Auxiliary
completed
 node_B_2
 Home Port: e5a
 172.17.26.12 172.17.26.11 DR Partner
completed
 Home Port: e5a
 172.17.26.12 172.17.26.10 DR Auxiliary
completed
 Home Port: e5b
 172.17.27.12 172.17.27.11 DR Partner
completed
 Home Port: e5b
 172.17.27.12 172.17.27.10 DR Auxiliary
completed
12 entries were displayed.
```

### Systems using shared MetroCluster/HA ports

The following output shows the IP addresses for a configuration with AFF A700 and FAS9000 systems with the MetroCluster IP interfaces on ports e5a and e5b. The interfaces can vary depending on the platform type.

```
cluster_B::*> metrocluster configuration-settings connection show
DR Source Destination
```

```

DR
Group Cluster Node Source Destination Partner Type
Config State Network Address Network Address

1 cluster_B
node_B_1
Home Port: e5a
172.17.26.13 172.17.26.12 HA Partner
completed
Home Port: e5a
172.17.26.13 172.17.26.10 DR Partner
completed
Home Port: e5a
172.17.26.13 172.17.26.11 DR Auxiliary
completed
Home Port: e5b
172.17.27.13 172.17.27.12 HA Partner
completed
Home Port: e5b
172.17.27.13 172.17.27.10 DR Partner
completed
Home Port: e5b
172.17.27.13 172.17.27.11 DR Auxiliary
completed
node_B_2
Home Port: e5a
172.17.26.12 172.17.26.13 HA Partner
completed
Home Port: e5a
172.17.26.12 172.17.26.11 DR Partner
completed
Home Port: e5a
172.17.26.12 172.17.26.10 DR Auxiliary
completed
Home Port: e5b
172.17.27.12 172.17.27.13 HA Partner
completed
Home Port: e5b
172.17.27.12 172.17.27.11 DR Partner
completed
Home Port: e5b
172.17.27.12 172.17.27.10 DR Auxiliary
completed
12 entries were displayed.

```

2. If you need to determine the VLAN ID or gateway address for the interface, determine the VLAN IDs from the surviving site:

```
metrocluster configuration-settings interface show
```

- You need to determine the VLAN ID if the platform models support VLAN IDs (see the [list above](#)) and if you are not using the default VLAN IDs.
- You need the gateway address if you are using [Layer 3 wide-area networks](#).

The VLAN IDs are included in the **Network Address** column of the output. The **Gateway** column shows the gateway IP address.

In this example the interfaces are e0a with the VLAN ID 120 and e0b with the VLAN ID 130:

```
Cluster-A::*> metrocluster configuration-settings interface show
DR
Config
Group Cluster Node Network Address Netmask Gateway
State

1
 cluster_A
 node_A_1
 Home Port: e0a-120
 172.17.26.10 255.255.255.0 -
completed
 Home Port: e0b-130
 172.17.27.10 255.255.255.0 -
completed
```

3. At the `LOADER` prompt for each of the disaster site nodes, set the `bootarg` value depending on whether your platform model uses shared cluster/HA ports or shared MetroCluster/HA ports:



- If the interfaces are using the default VLANs, or the platform model does not use a VLAN ID (see the [list above](#)), the `vlan-id` is not necessary.
- If the configuration is not using [Layer3 wide-area networks](#), the value for `gateway-IP-address` is **0** (zero).

### Systems using shared cluster/HA ports

Set the following bootarg:

```
setenv bootarg.mcc.port_a_ip_config local-IP-address/local-IP-
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id

setenv bootarg.mcc.port_b_ip_config local-IP-address/local-IP-
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id
```

The following commands set the values for node\_A\_1 using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12,120

setenv bootarg.mcc.port_b_ip_config
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12,130
```

The following example shows the commands for node\_A\_1 without a VLAN ID:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12

setenv bootarg.mcc.port_b_ip_config
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12
```

### Systems using shared MetroCluster/HA ports

Set the following bootarg:

```
setenv bootarg.mcc.port_a_ip_config local-IP-address/local-IP-
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-
address,vlan-id

setenv bootarg.mcc.port_b_ip_config local-IP-address/local-IP-
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-
address,vlan-id
```

The following commands set the values for node\_A\_1 using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12,120

setenv bootarg.mcc.port_b_ip_config
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12,130
```

The following example shows the commands for node\_A\_1 without a VLAN ID:

```
setenv bootarg.mcc.port_a_ip_config
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12

setenv bootarg.mcc.port_b_ip_config
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12
```

4. From the surviving site, gather the UUIDs for the disaster site:

```
metrocluster node show -fields node-cluster-uuid, node-uuid
```

```

cluster_B::> metrocluster node show -fields node-cluster-uuid, node-uuid

(metrocluster node show)
dr-group-id cluster node node-uuid
node-cluster-uuid

1 cluster_A node_A_1 f03cb63c-9a7e-11e7-b68b-00a098908039
ee7db9d5-9a82-11e7-b68b-00a098
908039
1 cluster_A node_A_2 aa9a7a7a-9a81-11e7-a4e9-00a098908c35
ee7db9d5-9a82-11e7-b68b-00a098
908039
1 cluster_B node_B_1 f37b240b-9ac1-11e7-9b42-00a098c9e55d
07958819-9ac6-11e7-9b42-00a098
c9e55d
1 cluster_B node_B_2 bf8e3f8f-9ac4-11e7-bd4e-00a098ca379f
07958819-9ac6-11e7-9b42-00a098
c9e55d
4 entries were displayed.
cluster_A::~*>

```

| Node      | UUID                                 |
|-----------|--------------------------------------|
| cluster_B | 07958819-9ac6-11e7-9b42-00a098c9e55d |
| node_B_1  | f37b240b-9ac1-11e7-9b42-00a098c9e55d |
| node_B_2  | bf8e3f8f-9ac4-11e7-bd4e-00a098ca379f |
| cluster_A | ee7db9d5-9a82-11e7-b68b-00a098908039 |
| node_A_1  | f03cb63c-9a7e-11e7-b68b-00a098908039 |
| node_A_2  | aa9a7a7a-9a81-11e7-a4e9-00a098908c35 |

5. At the replacement nodes' LOADER prompt, set the UUIDs:

```
setenv bootarg.mgwd.partner_cluster_uuid partner-cluster-UUID

setenv bootarg.mgwd.cluster_uuid local-cluster-UUID

setenv bootarg.mcc.pri_partner_uuid DR-partner-node-UUID

setenv bootarg.mcc.aux_partner_uuid DR-aux-partner-node-UUID

setenv bootarg.mcc_iscsi.node_uuid local-node-UUID`
```

a. Set the UUIDs on node\_A\_1.

The following example shows the commands for setting the UUIDs on node\_A\_1:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039

setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d

setenv bootarg.mcc.pri_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d

setenv bootarg.mcc.aux_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f

setenv bootarg.mcc_iscsi.node_uuid f03cb63c-9a7e-11e7-b68b-
00a098908039
```

b. Set the UUIDs on node\_A\_2:

The following example shows the commands for setting the UUIDs on node\_A\_2:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039

setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-
00a098c9e55d

setenv bootarg.mcc.pri_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-
00a098ca379f

setenv bootarg.mcc.aux_partner_uuid f37b240b-9ac1-11e7-9b42-
00a098c9e55d

setenv bootarg.mcc.iscsi.node_uuid aa9a7a7a-9a81-11e7-a4e9-
00a098908c35
```

6. If the original systems were configured for ADP, at each of the replacement nodes' LOADER prompt, enable ADP:

```
setenv bootarg.mcc.adp_enabled true
```

7. If running ONTAP 9.5, 9.6 or 9.7, at each of the replacement nodes' LOADER prompt, enable the following variable:

```
setenv bootarg.mcc.lun_part true
```

- a. Set the variables on node\_A\_1.

The following example shows the commands for setting the values on node\_A\_1 when running ONTAP 9.6:

```
setenv bootarg.mcc.lun_part true
```

- b. Set the variables on node\_A\_2.

The following example shows the commands for setting the values on node\_A\_2 when running ONTAP 9.6:

```
setenv bootarg.mcc.lun_part true
```

8. If the original systems were configured for end-to-end encryption, at each of the replacement nodes' LOADER prompt, set the following bootarg:

```
setenv bootarg.mccip.encryption_enabled 1
```

9. If the original systems were configured for ADP, at each of the replacement nodes' LOADER prompt, set the original system ID (**not** the system ID of the replacement controller module) and the system ID of the DR partner of the node:

```
setenv bootarg.mcc.local_config_id original-sysID
```

```
setenv bootarg.mcc.dr_partner dr_partner-sysID
```

### Determine the system IDs of the old controller modules

- a. Set the variables on node\_A\_1.

The following example shows the commands for setting the system IDs on node\_A\_1:

- The old system ID of node\_A\_1 is 4068741258.
- The system ID of node\_B\_1 is 4068741254.

```
setenv bootarg.mcc.local_config_id 4068741258
setenv bootarg.mcc.dr_partner 4068741254
```

- b. Set the variables on node\_A\_2.

The following example shows the commands for setting the system IDs on node\_A\_2:

- The old system ID of node\_A\_1 is 4068741260.
- The system ID of node\_B\_1 is 4068741256.

```
setenv bootarg.mcc.local_config_id 4068741260
setenv bootarg.mcc.dr_partner 4068741256
```

### Powering on the equipment at the disaster site (MetroCluster IP configurations)

You must power on the disk shelves and MetroCluster IP switches components at the disaster site. The controller modules at the disaster site remain at the LOADER prompt.

#### About this task

The examples in this procedure assume the following:

- Site A is the disaster site.
- Site B is the surviving site.

#### Steps

1. Turn on the disk shelves at the disaster site and make sure that all disks are running.
2. Turn on the MetroCluster IP switches if they are not already on.

### Configuring the IP switches (MetroCluster IP configurations)

You must configure any IP switches that were replaced.

#### About this task

This task applies to MetroCluster IP configurations only.

This must be done on both switches. Verify after configuring the first switch that storage access on the surviving site is not impacted.



You must not proceed with the second switch if storage access on the surviving site is impacted.

### Steps

1. Refer to [MetroCluster IP installation and configuration: : Differences among the ONTAP MetroCluster configurations](#) for procedures for cabling and configuring a replacement switch.

You can use the procedures in the following sections:

- Cabling the IP switches
  - Configuring the IP switches
2. If the ISLs were disabled at the surviving site, enable the ISLs and verify that the ISLs are online.
    - a. Enable the ISL interfaces on the first switch:

```
no shutdown
```

The following examples show the commands for a Broadcom IP switch or a Cisco IP switch.

| Switch vendor | Commands                                                                                                                                                                                                                                       |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Broadcom      | <pre>(IP_Switch_A_1)&gt; enable (IP_switch_A_1)# configure (IP_switch_A_1) (Config)# interface 0/13-0/16 (IP_switch_A_1) (Interface 0/13-0/16 )# no shutdown (IP_switch_A_1) (Interface 0/13-0/16 )# exit (IP_switch_A_1) (Config)# exit</pre> |
| Cisco         | <pre>IP_switch_A_1# conf t IP_switch_A_1(config)# int eth1/15-eth1/20 IP_switch_A_1(config)# no shutdown IP_switch_A_1(config)# copy running startup IP_switch_A_1(config)# show interface brief</pre>                                         |

- b. Enable the ISL interfaces on the partner switch:

```
no shutdown
```

The following examples show the commands for a Broadcom IP switch or a Cisco IP switch.

| Switch vendor | Commands                                                                                                                                                                                                                                       |
|---------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Broadcom      | <pre>(IP_Switch_A_2)&gt; enable (IP_switch_A_2)# configure (IP_switch_A_2) (Config)# interface 0/13-0/16 (IP_switch_A_2) (Interface 0/13-0/16 )# no shutdown (IP_switch_A_2) (Interface 0/13-0/16 )# exit (IP_switch_A_2) (Config)# exit</pre> |
| Cisco         | <pre>IP_switch_A_2# conf t IP_switch_A_2(config)# int eth1/15-eth1/20 IP_switch_A_2(config)# no shutdown IP_switch_A_2(config)# copy running startup IP_switch_A_2(config)# show interface brief</pre>                                         |

c. Verify that the interfaces are enabled:

```
show interface brief
```

The following example shows the output for a Cisco switch.

```

IP_switch_A_2(config)# show interface brief

Port VRF Status IP Address Speed MTU

mt0 -- up 10.10.99.10 100 1500

Ethernet VLAN Type Mode Status Reason Speed Port
Interface
#

.
.
.
Eth1/15 10 eth access up none 40G(D) --
Eth1/16 10 eth access up none 40G(D) --
Eth1/17 10 eth access down none auto(D) --
Eth1/18 10 eth access down none auto(D) --
Eth1/19 10 eth access down none auto(D) --
Eth1/20 10 eth access down none auto(D) --
.
.
.
IP_switch_A_2#

```

### Verify storage connectivity to the remote site (MetroCluster IP configurations)

You must confirm that the replaced nodes have connectivity to the disk shelves at the surviving site.

#### About this task

This task is performed on the replacement nodes at the disaster site.

This task is performed in Maintenance mode.

#### Steps

1. Display the disks that are owned by the original system ID.

```
disk show -s old-system-ID
```

The remote disks can be recognized by the 0m device. 0m indicates that the disk is connected via the MetroCluster iSCSI connection. These disks must be reassigned later in the recovery procedure.

```

*> disk show -s 4068741256
Local System ID: 1574774970

 DISK OWNER POOL SERIAL NUMBER HOME
DR HOME

0m.i0.0L11 node_A_2 (4068741256) Pool1 S396NA0HA02128 node_A_2
(4068741256) node_A_2 (4068741256)
0m.i0.1L38 node_A_2 (4068741256) Pool1 S396NA0J148778 node_A_2
(4068741256) node_A_2 (4068741256)
0m.i0.0L52 node_A_2 (4068741256) Pool1 S396NA0J148777 node_A_2
(4068741256) node_A_2 (4068741256)
...
...
NOTE: Currently 49 disks are unowned. Use 'disk show -n' for additional
information.
*>

```

2. Repeat this step on the other replacement nodes

### Reassigning disk ownership for pool 1 disks on the disaster site (MetroCluster IP configurations)

If one or both of the controller modules or NVRAM cards were replaced at the disaster site, the system ID has changed and you must reassign disks belonging to the root aggregates to the replacement controller modules.

#### About this task

Because the nodes are in switchover mode, only the disks containing the root aggregates of pool1 of the disaster site will be reassigned in this task. They are the only disks still owned by the old system ID at this point.

This task is performed on the replacement nodes at the disaster site.

This task is performed in Maintenance mode.

The examples make the following assumptions:

- Site A is the disaster site.
- node\_A\_1 has been replaced.
- node\_A\_2 has been replaced.
- Site B is the surviving site.
- node\_B\_1 is healthy.
- node\_B\_2 is healthy.

The old and new system IDs were identified in [Replace hardware and boot new controllers](#).

The examples in this procedure use controllers with the following system IDs:

| Node     | Original system ID | New system ID |
|----------|--------------------|---------------|
| node_A_1 | 4068741258         | 1574774970    |
| node_A_2 | 4068741260         | 1574774991    |
| node_B_1 | 4068741254         | unchanged     |
| node_B_2 | 4068741256         | unchanged     |

### Steps

1. With the replacement node in Maintenance mode, reassign the root aggregate disks, using the correct command, depending on whether your system is configured with ADP and your ONTAP version.

You can proceed with the reassignment when prompted.

| If the system is using ADP... | Use this command for disk reassignment...                                             |
|-------------------------------|---------------------------------------------------------------------------------------|
| Yes (ONTAP 9.8)               | <code>disk reassign -s old-system-ID -d new-system-ID -r dr-partner-system-ID</code>  |
| Yes (ONTAP 9.7.x and earlier) | <code>disk reassign -s old-system-ID -d new-system-ID -p old-partner-system-ID</code> |
| No                            | <code>disk reassign -s old-system-ID -d new-system-ID</code>                          |

The following example shows reassignment of drives on a non-ADP system:

```

*> disk reassign -s 4068741256 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from
maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode.
Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and
giveback of the HA partner node to ensure disk reassignment is
successful.
Do you want to continue (y/n)? y
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 537037643.
Do you want to continue (y/n)? y
disk reassign parameters: new_home_owner_id 537070473 ,
new_home_owner_name
Disk 0m.i0.3L14 will be reassigned.
Disk 0m.i0.1L6 will be reassigned.
Disk 0m.i0.1L8 will be reassigned.
Number of disks to be reassigned: 3

```

## 2. Destroy the contents of the mailbox disks:

```
mailbox destroy local
```

You can proceed with the destroy operation when prompted.

The following example shows the output for the mailbox destroy local command:

```

*> mailbox destroy local
Destroying mailboxes forces a node to create new empty mailboxes,
which clears any takeover state, removes all knowledge
of out-of-date plexes of mirrored volumes, and will prevent
management services from going online in 2-node cluster
HA configurations.
Are you sure you want to destroy the local mailboxes? y
.....Mailboxes destroyed.
*>

```

## 3. If disks have been replaced, there will be failed local plexes that must be deleted.

### a. Display the aggregate status:

```
aggr status
```

In the following example, plex node\_A\_1\_aggr0/plex0 has failed.

```

*> aggr status
Aug 18 15:00:07 [node_B_1:raid.vol.mirror.degraded:ALERT]: Aggregate
node_A_1_aggr0 is
 mirrored and one plex has failed. It is no longer protected by
 mirroring.
Aug 18 15:00:07 [node_B_1:raid.debug:info]: Mirrored aggregate
node_A_1_aggr0 has plex0
 clean(-1), online(0)
Aug 18 15:00:07 [node_B_1:raid.debug:info]: Mirrored aggregate
node_A_1_aggr0 has plex2
 clean(0), online(1)
Aug 18 15:00:07 [node_B_1:raid.mirror.vote.noRecord1Plex:error]:
WARNING: Only one plex
 in aggregate node_A_1_aggr0 is available. Aggregate might contain
 stale data.
Aug 18 15:00:07 [node_B_1:raid.debug:info]:
volobj_mark_sb_recovery_aggrs: tree:
 node_A_1_aggr0 vol_state:1 mcc_dr_opstate: unknown
Aug 18 15:00:07 [node_B_1:raid.fsm.commitStateTransit:debug]:
/node_A_1_aggr0 (VOL):
 raid state change UNINITD -> NORMAL
Aug 18 15:00:07 [node_B_1:raid.fsm.commitStateTransit:debug]:
/node_A_1_aggr0 (MIRROR):
 raid state change UNINITD -> DEGRADED
Aug 18 15:00:07 [node_B_1:raid.fsm.commitStateTransit:debug]:
/node_A_1_aggr0/plex0
 (PLEX): raid state change UNINITD -> FAILED
Aug 18 15:00:07 [node_B_1:raid.fsm.commitStateTransit:debug]:
/node_A_1_aggr0/plex2
 (PLEX): raid state change UNINITD -> NORMAL
Aug 18 15:00:07 [node_B_1:raid.fsm.commitStateTransit:debug]:
/node_A_1_aggr0/plex2/rg0
 (GROUP): raid state change UNINITD -> NORMAL
Aug 18 15:00:07 [node_B_1:raid.debug:info]: Topology updated for
aggregate node_A_1_aggr0
 to plex plex2
*>

```

**b. Delete the failed plex:**

```
aggr destroy plex-id
```

```
*> aggr destroy node_A_1_aggr0/plex0
```

4. Halt the node to display the LOADER prompt:

```
halt
```

5. Repeat these steps on the other node at the disaster site.

## Booting to ONTAP on replacement controller modules in MetroCluster IP configurations

You must boot the replacement nodes at the disaster site to the ONTAP operating system.

### About this task

This task begins with the nodes at the disaster site in Maintenance mode.

### Steps

1. On one of the replacement nodes, exit to the LOADER prompt: `halt`
2. Display the boot menu: `boot_ontap menu`
3. From the boot menu, select option 6, **Update flash from backup config**.

The system boots twice. You should respond `yes` when prompted to continue. After the second boot, you should respond `y` when prompted about the system ID mismatch.



If you did not clear the NVRAM contents of a used replacement controller module, then you might see the following panic message: `PANIC: NVRAM contents are invalid...` If this occurs, boot the system to the ONTAP prompt again (`boot_ontap menu`). You then need to [Reset the boot\\_recovery and rdb\\_corrupt bootargs](#)

- Confirmation to continue prompt:

```
Selection (1-9)? 6
```

```
This will replace all flash-based configuration with the last backup
to
disks. Are you sure you want to continue?: yes
```

- System ID mismatch prompt:

```
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
```

4. From the surviving site, verify that the correct partner system IDs have been applied to the nodes:

```
metrocluster node show -fields node-systemid,ha-partner-systemid,dr-partner-
systemid,dr-auxiliary-systemid
```

In this example, the following new system IDs should appear in the output:

- Node\_A\_1: 1574774970
- Node\_A\_2: 1574774991

The "ha-partner-systemid" column should show the new system IDs.

```
metrocluster node show -fields node-systemid,ha-partner-systemid,dr-
partner-systemid,dr-auxiliary-systemid

dr-group-id cluster node node-systemid ha-partner-systemid dr-
partner-systemid dr-auxiliary-systemid

1 Cluster_A Node_A_1 1574774970 1574774991
4068741254 4068741256
1 Cluster_A Node_A_2 1574774991 1574774970
4068741256 4068741254
1 Cluster_B Node_B_1 - - -
-
1 Cluster_B Node_B_2 - - -
-
4 entries were displayed.
```

5. If the partner system IDs were not correctly set, you must manually set the correct value:
  - a. Halt and display the LOADER prompt on the node.
  - b. Verify the partner-sysID bootarg's current value:

```
printenv
```

- c. Set the value to the correct partner system ID:

```
setenv partner-sysid partner-sysID
```

- d. Boot the node:

```
boot_ontap
```

- e. Repeat these substeps on the other node, if necessary.

6. Confirm that the replacement nodes at the disaster site are ready for switchback:

```
metrocluster node show
```

The replacement nodes should be in waiting for switchback recovery mode. If they are in normal mode instead, you can reboot the replacement nodes. After that boot, the nodes should be in waiting for switchback recovery mode.

The following example shows that the replacement nodes are ready for switchback:

```

cluster_B::> metrocluster node show
DR
Group Cluster Node Configuration DR
State Mirroring Mode

1 cluster_B
 node_B_1 configured enabled switchover
completed
 node_B_2 configured enabled switchover
completed
 cluster_A
 node_A_1 configured enabled waiting for
switchback recovery
 node_A_2 configured enabled waiting for
switchback recovery
4 entries were displayed.

cluster_B::>

```

#### 7. Verify the MetroCluster connection configuration settings:

```
metrocluster configuration-settings connection show
```

The configuration state should indicate completed.

```

cluster_B::*> metrocluster configuration-settings connection show
DR
Group Cluster Node Source Destination
Config State Network Address Network Address Partner Type

1 cluster_B
 node_B_2
 Home Port: e5a
 172.17.26.13 172.17.26.12 HA Partner
completed
 Home Port: e5a
 172.17.26.13 172.17.26.10 DR Partner
completed
 Home Port: e5a
 172.17.26.13 172.17.26.11 DR Auxiliary
completed
 Home Port: e5b
 172.17.27.13 172.17.27.12 HA Partner
completed

```

```

 Home Port: e5b
 172.17.27.13 172.17.27.10 DR Partner
completed

 Home Port: e5b
 172.17.27.13 172.17.27.11 DR Auxiliary
completed
node_B_1
 Home Port: e5a
 172.17.26.12 172.17.26.13 HA Partner
completed

 Home Port: e5a
 172.17.26.12 172.17.26.11 DR Partner
completed

 Home Port: e5a
 172.17.26.12 172.17.26.10 DR Auxiliary
completed

 Home Port: e5b
 172.17.27.12 172.17.27.13 HA Partner
completed

 Home Port: e5b
 172.17.27.12 172.17.27.11 DR Partner
completed

 Home Port: e5b
 172.17.27.12 172.17.27.10 DR Auxiliary
completed
cluster_A
node_A_2
 Home Port: e5a
 172.17.26.11 172.17.26.10 HA Partner
completed

 Home Port: e5a
 172.17.26.11 172.17.26.12 DR Partner
completed

 Home Port: e5a
 172.17.26.11 172.17.26.13 DR Auxiliary
completed

 Home Port: e5b
 172.17.27.11 172.17.27.10 HA Partner
completed

 Home Port: e5b
 172.17.27.11 172.17.27.12 DR Partner
completed

 Home Port: e5b
 172.17.27.11 172.17.27.13 DR Auxiliary
completed
node_A_1

```

```

Home Port: e5a
172.17.26.10 172.17.26.11 HA Partner
completed
Home Port: e5a
172.17.26.10 172.17.26.13 DR Partner
completed
Home Port: e5a
172.17.26.10 172.17.26.12 DR Auxiliary
completed
Home Port: e5b
172.17.27.10 172.17.27.11 HA Partner
completed
Home Port: e5b
172.17.27.10 172.17.27.13 DR Partner
completed
Home Port: e5b
172.17.27.10 172.17.27.12 DR Auxiliary
completed
24 entries were displayed.

cluster_B::*>

```

8. Repeat the previous steps on the other node at the disaster site.

### Reset the boot\_recovery and rdb\_corrupt bootargs

If required, you can reset the boot\_recovery and rdb\_corrupt\_bootargs

#### Steps

1. Halt the node back to the LOADER prompt:

```
siteA::*> halt -node <node-name>
```

2. Check if the following bootargs have been set:

```
LOADER> printenv bootarg.init.boot_recovery
LOADER> printenv bootarg.rdb_corrupt
```

3. If either bootarg has been set to a value, unset it and boot ONTAP:

```
LOADER> unsetenv bootarg.init.boot_recovery
LOADER> unsetenv bootarg.rdb_corrupt
LOADER> saveenv
LOADER> bye
```

## Restoring connectivity from the surviving nodes to the disaster site (MetroCluster IP configurations)

You must restore the MetroCluster iSCSI initiator connections from the surviving nodes.

### About this task

This procedure is only required on MetroCluster IP configurations.

### Steps

1. From either surviving node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).

2. Connect the iSCSI initiators on both surviving nodes in the DR group:

```
storage iscsi-initiator connect -node surviving-node -label *
```

The following example shows the commands for connecting the initiators on site B:

```
site_B::*> storage iscsi-initiator connect -node node_B_1 -label *
site_B::*> storage iscsi-initiator connect -node node_B_2 -label *
```

3. Return to the admin privilege level:

```
set -privilege admin
```

## Verifying automatic assignment or manually assigning pool 0 drives

On systems configured for ADP, you must verify that pool 0 drives have been automatically assigned. On systems that are not configured for ADP, you must manually assign the pool 0 drives.

### Verifying drive assignment of pool 0 drives on ADP systems at the disaster site (MetroCluster IP systems)

If drives have been replaced at the disaster site and the system is configured for ADP, you must verify that the remote drives are visible to the nodes and have been assigned correctly.

### Step

1. Verify that pool 0 drives are assigned automatically:

disk show

In the following example for an AFF A800 system with no external shelves, one quarter (8 drives) were automatically assigned to node\_A\_1 and one quarter were automatically assigned to node\_A\_2. The remaining drives will be remote (pool1) drives for node\_B\_1 and node\_B\_2.

```
cluster_A::*> disk show
```

| Disk Owner       | Usable Size | Disk Shelf | Bay | Container Type | Type       | Container Name |
|------------------|-------------|------------|-----|----------------|------------|----------------|
| node_A_1:0n.12   | 1.75TB      | 0          | 12  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.13   | 1.75TB      | 0          | 13  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.14   | 1.75TB      | 0          | 14  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.15   | 1.75TB      | 0          | 15  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.16   | 1.75TB      | 0          | 16  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.17   | 1.75TB      | 0          | 17  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.18   | 1.75TB      | 0          | 18  | SSD-NVM        | shared     | aggr0          |
| node_A_1         |             |            |     |                |            |                |
| node_A_1:0n.19   | 1.75TB      | 0          | 19  | SSD-NVM        | shared     | -              |
| node_A_1         |             |            |     |                |            |                |
| node_A_2:0n.0    | 1.75TB      | 0          | 0   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.1    | 1.75TB      | 0          | 1   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.2    | 1.75TB      | 0          | 2   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.3    | 1.75TB      | 0          | 3   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.4    | 1.75TB      | 0          | 4   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.5    | 1.75TB      | 0          | 5   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.6    | 1.75TB      | 0          | 6   | SSD-NVM        | shared     |                |
| aggr0_node_A_2_0 | node_A_2    |            |     |                |            |                |
| node_A_2:0n.7    | 1.75TB      | 0          | 7   | SSD-NVM        | shared     | -              |
| node_A_2         |             |            |     |                |            |                |
| node_A_2:0n.24   | -           | 0          | 24  | SSD-NVM        | unassigned | -              |
| node_A_2:0n.25   | -           | 0          | 25  | SSD-NVM        | unassigned | -              |

```

node_A_2:0n.26 - 0 26 SSD-NVM unassigned - -
node_A_2:0n.27 - 0 27 SSD-NVM unassigned - -
node_A_2:0n.28 - 0 28 SSD-NVM unassigned - -
node_A_2:0n.29 - 0 29 SSD-NVM unassigned - -
node_A_2:0n.30 - 0 30 SSD-NVM unassigned - -
node_A_2:0n.31 - 0 31 SSD-NVM unassigned - -
node_A_2:0n.36 - 0 36 SSD-NVM unassigned - -
node_A_2:0n.37 - 0 37 SSD-NVM unassigned - -
node_A_2:0n.38 - 0 38 SSD-NVM unassigned - -
node_A_2:0n.39 - 0 39 SSD-NVM unassigned - -
node_A_2:0n.40 - 0 40 SSD-NVM unassigned - -
node_A_2:0n.41 - 0 41 SSD-NVM unassigned - -
node_A_2:0n.42 - 0 42 SSD-NVM unassigned - -
node_A_2:0n.43 - 0 43 SSD-NVM unassigned - -
32 entries were displayed.

```

### Assigning pool 0 drives on non-ADP systems at the disaster site (MetroCluster IP configurations)

If drives have been replaced at the disaster site and the system is not configured for ADP, you need to manually assign new drives to pool 0.

#### About this task

For ADP systems, the drives are assigned automatically.

#### Steps

1. On one of the replacement nodes at the disaster site, reassign the node's pool 0 drives:

```
storage disk assign -n number-of-replacement disks -p 0
```

This command assigns the newly added (and unowned) drives on the disaster site. You should assign the same number and size (or larger) of drives that the node had prior to the disaster. The `storage disk assign` man page contains more information about performing more granular drive assignment.

2. Repeat the step on the other replacement node at the disaster site.

### Assigning pool 1 drives on the surviving site (MetroCluster IP configurations)

If drives have been replaced at the disaster site and the system is not configured for ADP, at the surviving site you need to manually assign remote drives located at the disaster site to the surviving nodes' pool 1. You must identify the number of drives to assign.

#### About this task

For ADP systems, the drives are assigned automatically.

#### Step

1. On the surviving site, assign the first node's pool 1 (remote) drives: `storage disk assign -n number-of-replacement disks -p 1 0m*`

This command assigns the newly added and unowned drives on the disaster site.

The following command assigns 22 drives:

```
cluster_B::> storage disk assign -n 22 -p 1 0m*
```

### Deleting failed plexes owned by the surviving site (MetroCluster IP configurations)

After replacing hardware and assigning disks, you must delete failed remote plexes that are owned by the surviving site nodes but located at the disaster site.

#### About this task

These steps are performed on the surviving cluster.

#### Steps

1. Identify the local aggregates:

```
storage aggregate show -is-home true
```

```
cluster_B::> storage aggregate show -is-home true

cluster_B Aggregates:
Aggregate Size Available Used% State #Vols Nodes RAID
Status

node_B_1_aggr0 1.49TB 74.12GB 95% online 1 node_B_1
raid4,

mirror

degraded
node_B_2_aggr0 1.49TB 74.12GB 95% online 1 node_B_2
raid4,

mirror

degraded
node_B_1_aggr1 2.99TB 2.88TB 3% online 15 node_B_1
raid_dp,

mirror

degraded
node_B_1_aggr2 2.99TB 2.91TB 3% online 14 node_B_1
raid_tec,

mirror
```

```
degraded
node_B_2_aggr1 2.95TB 2.80TB 5% online 37 node_B_2
raid_dp,

mirror

degraded
node_B_2_aggr2 2.99TB 2.87TB 4% online 35 node_B_2
raid_tec,

mirror

degraded
6 entries were displayed.

cluster_B::>
```

2. Identify the failed remote plexes:

```
storage aggregate plex show
```

The following example calls out the plexes that are remote (not plex0) and have a status of "failed":

```

cluster_B::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate plex status is-online pool

node_B_1_aggr0 plex0 normal,active true 0
node_B_1_aggr0 plex4 failed,inactive false - <<<<---Plex at remote site
node_B_2_aggr0 plex0 normal,active true 0
node_B_2_aggr0 plex4 failed,inactive false - <<<<---Plex at remote site
node_B_1_aggr1 plex0 normal,active true 0
node_B_1_aggr1 plex4 failed,inactive false - <<<<---Plex at remote site
node_B_1_aggr2 plex0 normal,active true 0
node_B_1_aggr2 plex1 failed,inactive false - <<<<---Plex at remote site
node_B_2_aggr1 plex0 normal,active true 0
node_B_2_aggr1 plex4 failed,inactive false - <<<<---Plex at remote site
node_B_2_aggr2 plex0 normal,active true 0
node_B_2_aggr2 plex1 failed,inactive false - <<<<---Plex at remote site
node_A_1_aggr1 plex0 failed,inactive false -
node_A_1_aggr1 plex4 normal,active true 1
node_A_1_aggr2 plex0 failed,inactive false -
node_A_1_aggr2 plex1 normal,active true 1
node_A_2_aggr1 plex0 failed,inactive false -
node_A_2_aggr1 plex4 normal,active true 1
node_A_2_aggr2 plex0 failed,inactive false -
node_A_2_aggr2 plex1 normal,active true 1
20 entries were displayed.

cluster_B::>

```

3. Take offline each of the failed plexes, and then delete them:

a. Take offline the failed plexes:

```
storage aggregate plex offline -aggregate aggregate-name -plex plex-id
```

The following example shows the aggregate "node\_B\_2\_aggr1/plex1" being taken offline:

```

cluster_B::> storage aggregate plex offline -aggregate node_B_1_aggr0
-plex plex4

Plex offline successful on plex: node_B_1_aggr0/plex4

```

b. Delete the failed plex:

```
storage aggregate plex delete -aggregate aggregate-name -plex plex-id
```

You can destroy the plex when prompted.

The following example shows the plex node\_B\_2\_aggr1/plex1 being deleted.

```
cluster_B::> storage aggregate plex delete -aggregate node_B_1_aggr0
-plex plex4

Warning: Aggregate "node_B_1_aggr0" is being used for the local
management root
 volume or HA partner management root volume, or has been
marked as
 the aggregate to be used for the management root volume
after a
 reboot operation. Deleting plex "plex4" for this aggregate
could lead
 to unavailability of the root volume after a disaster
recovery
 procedure. Use the "storage aggregate show -fields
 has-mroot,has-partner-mroot,root" command to view such
aggregates.

Warning: Deleting plex "plex4" of mirrored aggregate "node_B_1_aggr0"
on node
 "node_B_1" in a MetroCluster configuration will disable its
synchronous disaster recovery protection. Are you sure you
want to
 destroy this plex? {y|n}: y
[Job 633] Job succeeded: DONE

cluster_B::>
```

You must repeat these steps for each of the failed plexes.

#### 4. Confirm that the plexes have been removed:

```
storage aggregate plex show -fields aggregate,status,is-online,plex,pool
```

```

cluster_B::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate plex status is-online pool

node_B_1_aggr0 plex0 normal,active true 0
node_B_2_aggr0 plex0 normal,active true 0
node_B_1_aggr1 plex0 normal,active true 0
node_B_1_aggr2 plex0 normal,active true 0
node_B_2_aggr1 plex0 normal,active true 0
node_B_2_aggr2 plex0 normal,active true 0
node_A_1_aggr1 plex0 failed,inactive false -
node_A_1_aggr1 plex4 normal,active true 1
node_A_1_aggr2 plex0 failed,inactive false -
node_A_1_aggr2 plex1 normal,active true 1
node_A_2_aggr1 plex0 failed,inactive false -
node_A_2_aggr1 plex4 normal,active true 1
node_A_2_aggr2 plex0 failed,inactive false -
node_A_2_aggr2 plex1 normal,active true 1
14 entries were displayed.

cluster_B::>

```

5. Identify the switched-over aggregates:

```
storage aggregate show -is-home false
```

You can also use the `storage aggregate plex show -fields aggregate,status,is-online,plex,pool` command to identify plex 0 switched-over aggregates. They will have a status of "failed, inactive".

The following commands show four switched-over aggregates:

- `node_A_1_aggr1`
- `node_A_1_aggr2`
- `node_A_2_aggr1`
- `node_A_2_aggr2`

```

cluster_B::> storage aggregate show -is-home false

cluster_A Switched Over Aggregates:
Aggregate Size Available Used% State #Vols Nodes RAID
Status

node_A_1_aggr1 2.12TB 1.88TB 11% online 91 node_B_1
raid_dp,

mirror

degraded
node_A_1_aggr2 2.89TB 2.64TB 9% online 90 node_B_1
raid_tec,

mirror

degraded
node_A_2_aggr1 2.12TB 1.86TB 12% online 91 node_B_2
raid_dp,

mirror

degraded
node_A_2_aggr2 2.89TB 2.64TB 9% online 90 node_B_2
raid_tec,

mirror

degraded
4 entries were displayed.

cluster_B::>

```

#### 6. Identify switched-over plexes:

```
storage aggregate plex show -fields aggregate,status,is-online,Plex,pool
```

You want to identify the plexes with a status of "failed, inactive".

The following commands show four switched-over aggregates:

```

cluster_B::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate plex status is-online pool

node_B_1_aggr0 plex0 normal,active true 0
node_B_2_aggr0 plex0 normal,active true 0
node_B_1_aggr1 plex0 normal,active true 0
node_B_1_aggr2 plex0 normal,active true 0
node_B_2_aggr1 plex0 normal,active true 0
node_B_2_aggr2 plex0 normal,active true 0
node_A_1_aggr1 plex0 failed,inactive false - <<<<-- Switched over
aggr/Plex0
node_A_1_aggr1 plex4 normal,active true 1
node_A_1_aggr2 plex0 failed,inactive false - <<<<-- Switched over
aggr/Plex0
node_A_1_aggr2 plex1 normal,active true 1
node_A_2_aggr1 plex0 failed,inactive false - <<<<-- Switched over
aggr/Plex0
node_A_2_aggr1 plex4 normal,active true 1
node_A_2_aggr2 plex0 failed,inactive false - <<<<-- Switched over
aggr/Plex0
node_A_2_aggr2 plex1 normal,active true 1
14 entries were displayed.

cluster_B::>

```

## 7. Delete the failed plex:

```
storage aggregate plex delete -aggregate node_A_1_aggr1 -plex plex0
```

You can destroy the plex when prompted.

The following example shows the plex node\_A\_1\_aggr1/plex0 being deleted:

```
cluster_B::> storage aggregate plex delete -aggregate node_A_1_aggr1
-plex plex0

Warning: Aggregate "node_A_1_aggr1" hosts MetroCluster metadata volume
"MDV_CRS_e8457659b8a711e78b3b00a0988fe74b_A". Deleting plex
"plex0"
 for this aggregate can lead to the failure of configuration
 replication across the two DR sites. Use the "volume show
-vserver
 <admin-vserver> -volume MDV_CRS*" command to verify the
location of
 such volumes.

Warning: Deleting plex "plex0" of mirrored aggregate "node_A_1_aggr1" on
node
 "node_A_1" in a MetroCluster configuration will disable its
 synchronous disaster recovery protection. Are you sure you want
to
 destroy this plex? {y|n}: y
[Job 639] Job succeeded: DONE

cluster_B::>
```

You must repeat these steps for each of the failed aggregates.

8. Verify that there are no failed plexes remaining on the surviving site.

The following output shows that all plexes are normal, active, and online.

```

cluster_B::> storage aggregate plex show -fields aggregate,status,is-
online,Plex,pool
aggregate plex status is-online pool

node_B_1_aggr0 plex0 normal,active true 0
node_B_2_aggr0 plex0 normal,active true 0
node_B_1_aggr1 plex0 normal,active true 0
node_B_2_aggr2 plex0 normal,active true 0
node_B_1_aggr1 plex0 normal,active true 0
node_B_2_aggr2 plex0 normal,active true 0
node_A_1_aggr1 plex4 normal,active true 1
node_A_1_aggr2 plex1 normal,active true 1
node_A_2_aggr1 plex4 normal,active true 1
node_A_2_aggr2 plex1 normal,active true 1
10 entries were displayed.

cluster_B::>

```

### Performing aggregate healing and restoring mirrors (MetroCluster IP configurations)

After replacing hardware and assigning disks, in systems running ONTAP 9.5 or earlier you can perform the MetroCluster healing operations. In all versions of ONTAP, you must then confirm that aggregates are mirrored and, if necessary, restart mirroring.

#### About this task

Beginning with ONTAP 9.6, the healing operations are performed automatically when the disaster site nodes boot up. The healing commands are not required.

These steps are performed on the surviving cluster.

#### Steps

1. If you are using ONTAP 9.6 or later, you must verify that automatic healing completed successfully:
  - a. Confirm that the heal-aggr-auto and heal-root-aggr-auto operations completed:

```
metrocluster operation history show
```

The following output shows that the operations have completed successfully on cluster\_A.

```

cluster_B::*> metrocluster operation history show
Operation State Start Time End
Time

heal-root-aggr-auto successful 2/25/2019 06:45:58
2/25/2019 06:46:02
heal-aggr-auto successful 2/25/2019 06:45:48
2/25/2019 06:45:52
.
.
.

```

b. Confirm that the disaster site is ready for switchback:

```
metrocluster node show
```

The following output shows that the operations have completed successfully on cluster\_A.

```

cluster_B::*> metrocluster node show
DR Configuration DR
Group Cluster Node State Mirroring Mode

1 cluster_A
 node_A_1 configured enabled heal roots
completed
 node_A_2 configured enabled heal roots
completed
 cluster_B
 node_B_1 configured enabled waiting for
switchback recovery
 node_B_2 configured enabled waiting for
switchback recovery
4 entries were displayed.

```

2. If you are using ONTAP 9.5 or earlier, you must perform aggregate healing:

a. Verify the state of the nodes:

```
metrocluster node show
```

The following output shows that switchover has completed, so healing can be performed.

```

cluster_B::> metrocluster node show
DR Configuration DR
Group Cluster Node State Mirroring Mode

1 cluster_B
 node_B_1 configured enabled switchover
completed
 node_B_2 configured enabled switchover
completed
 cluster_A
 node_A_1 configured enabled waiting for
switchback recovery
 node_A_2 configured enabled waiting for
switchback recovery
4 entries were displayed.

cluster_B::>

```

b. Perform the aggregates healing phase:

```
metrocluster heal -phase aggregates
```

The following output shows a typical aggregates healing operation.

```

cluster_B::*> metrocluster heal -phase aggregates
[Job 647] Job succeeded: Heal Aggregates is successful.

cluster_B::*> metrocluster operation show
 Operation: heal-aggregates
 State: successful
 Start Time: 10/26/2017 12:01:15
 End Time: 10/26/2017 12:01:17
 Errors: -

cluster_B::*>

```

c. Verify that aggregate healing has completed and the disaster site is ready for switchback:

```
metrocluster node show
```

The following output shows that the "heal aggregates" phase has completed on cluster\_A.

```

cluster_B::> metrocluster node show
DR
Group Cluster Node Configuration State DR Mirroring Mode

1 cluster_A
node_A_1 configured enabled heal
aggregates completed
node_A_2 configured enabled heal
aggregates completed
cluster_B
node_B_1 configured enabled waiting for
switchback recovery
node_B_2 configured enabled waiting for
switchback recovery
4 entries were displayed.

cluster_B::>

```

3. If disks have been replaced, you must mirror the local and switched-over aggregates:

a. Display the aggregates:

```
storage aggregate show
```

```

cluster_B::> storage aggregate show
cluster_B Aggregates:
Aggregate Size Available Used% State #Vols Nodes
RAID Status

node_B_1_aggr0 1.49TB 74.12GB 95% online 1 node_B_1
raid4,
normal
node_B_2_aggr0 1.49TB 74.12GB 95% online 1 node_B_2
raid4,
normal
node_B_1_aggr1 3.14TB 3.04TB 3% online 15 node_B_1
raid_dp,
normal
node_B_1_aggr2 3.14TB 3.06TB 3% online 14 node_B_1
raid_tec,

```

```

normal
node_B_1_aggr1 3.14TB 2.99TB 5% online 37 node_B_2
raid_dp,

normal
node_B_1_aggr2 3.14TB 3.02TB 4% online 35 node_B_2
raid_tec,

normal

cluster_A Switched Over Aggregates:
Aggregate Size Available Used% State #Vols Nodes
RAID Status

node_A_1_aggr1 2.36TB 2.12TB 10% online 91 node_B_1
raid_dp,

normal
node_A_1_aggr2 3.14TB 2.90TB 8% online 90 node_B_1
raid_tec,

normal
node_A_2_aggr1 2.36TB 2.10TB 11% online 91 node_B_2
raid_dp,

normal
node_A_2_aggr2 3.14TB 2.89TB 8% online 90 node_B_2
raid_tec,

normal
12 entries were displayed.

cluster_B::>

```

**b. Mirror the aggregate:**

```
storage aggregate mirror -aggregate aggregate-name
```

The following output shows a typical mirroring operation.

```

cluster_B::> storage aggregate mirror -aggregate node_B_1_aggr1

Info: Disks would be added to aggregate "node_B_1_aggr1" on node
"node_B_1" in
 the following manner:

 Second Plex

 RAID Group rg0, 6 disks (block checksum, raid_dp)
 Position Disk Type
Size

- dparity 5.20.6 SSD
- parity 5.20.14 SSD
- data 5.21.1 SSD
894.0GB data 5.21.3 SSD
894.0GB data 5.22.3 SSD
894.0GB data 5.21.13 SSD
894.0GB

Aggregate capacity available for volume use would be 2.99TB.

Do you want to continue? {y|n}: y

```

- c. Repeat the previous step for each of the aggregates from the surviving site.
- d. Wait for the aggregates to resynchronize; you can check the status with the `storage aggregate show` command.

The following output shows that a number of aggregates are resynchronizing.

```

cluster_B::> storage aggregate show

cluster_B Aggregates:
Aggregate Size Available Used% State #Vols Nodes
RAID Status

node_B_1_aggr0 1.49TB 74.12GB 95% online 1 node_B_1
raid4,

```

```

mirrored,

normal
node_B_2_aggr0 1.49TB 74.12GB 95% online 1 node_B_2
raid4,

mirrored,

normal
node_B_1_aggr1 2.86TB 2.76TB 4% online 15 node_B_1
raid_dp,

resyncing
node_B_1_aggr2 2.89TB 2.81TB 3% online 14 node_B_1
raid_tec,

resyncing
node_B_2_aggr1 2.73TB 2.58TB 6% online 37 node_B_2
raid_dp,

resyncing
node_B-2_aggr2 2.83TB 2.71TB 4% online 35 node_B_2
raid_tec,

resyncing

cluster_A Switched Over Aggregates:
Aggregate Size Available Used% State #Vols Nodes
RAID Status

node_A_1_aggr1 1.86TB 1.62TB 13% online 91 node_B_1
raid_dp,

resyncing
node_A_1_aggr2 2.58TB 2.33TB 10% online 90 node_B_1
raid_tec,

resyncing
node_A_2_aggr1 1.79TB 1.53TB 14% online 91 node_B_2
raid_dp,

resyncing
node_A_2_aggr2 2.64TB 2.39TB 9% online 90 node_B_2
raid_tec,

```

```
resyncing
12 entries were displayed.
```

e. Confirm that all aggregates are online and have resynchronized:

```
storage aggregate plex show
```

The following output shows that all aggregates have resynchronized.

```
cluster_A::> storage aggregate plex show
()
Aggregate Plex Is Is Resyncing
 Online Resyncing Percent Status

node_B_1_aggr0 plex0 true false - normal,active
node_B_1_aggr0 plex8 true false - normal,active
node_B_2_aggr0 plex0 true false - normal,active
node_B_2_aggr0 plex8 true false - normal,active
node_B_1_aggr1 plex0 true false - normal,active
node_B_1_aggr1 plex9 true false - normal,active
node_B_1_aggr2 plex0 true false - normal,active
node_B_1_aggr2 plex5 true false - normal,active
node_B_2_aggr1 plex0 true false - normal,active
node_B_2_aggr1 plex9 true false - normal,active
node_B_2_aggr2 plex0 true false - normal,active
node_B_2_aggr2 plex5 true false - normal,active
node_A_1_aggr1 plex4 true false - normal,active
node_A_1_aggr1 plex8 true false - normal,active
node_A_1_aggr2 plex1 true false - normal,active
node_A_1_aggr2 plex5 true false - normal,active
node_A_2_aggr1 plex4 true false - normal,active
node_A_2_aggr1 plex8 true false - normal,active
node_A_2_aggr2 plex1 true false - normal,active
node_A_2_aggr2 plex5 true false - normal,active
20 entries were displayed.
```

4. On systems running ONTAP 9.5 and earlier, perform the root-aggregates healing phase:

```
metrocluster heal -phase root-aggregates
```

```

cluster_B::> metrocluster heal -phase root-aggregates
[Job 651] Job is queued: MetroCluster Heal Root Aggregates Job.Oct 26
13:05:00
[Job 651] Job succeeded: Heal Root Aggregates is successful.

```

5. Verify that the "heal roots" phase has completed and the disaster site is ready for switchback:

The following output shows that the "heal roots" phase has completed on cluster\_A.

```

cluster_B::> metrocluster node show
DR Configuration DR
Group Cluster Node State Mirroring Mode

1 cluster_A
 node_A_1 configured enabled heal roots
completed
 node_A_2 configured enabled heal roots
completed
 cluster_B
 node_B_1 configured enabled waiting for
switchback recovery
 node_B_2 configured enabled waiting for
switchback recovery
4 entries were displayed.

cluster_B::>

```

Proceed to verify the licenses on the replaced nodes.

[Verifying licenses on the replaced nodes](#)

## Prepare for switchback in a MetroCluster FC configuration

### Verifying port configuration (MetroCluster FC configurations only)

You must set the environmental variables on the node and then power it off to prepare it for MetroCluster configuration.

#### About this task

This procedure is performed with the replacement controller modules in Maintenance mode.

The steps to check configuration of ports is needed only on systems in which FC or CNA ports are used in initiator mode.

#### Steps

1. In Maintenance mode, restore the FC port configuration:

```
ucadmin modify -m fc -t initiatoradapter_name
```

If you only want to use one of a port pair in the initiator configuration, enter a precise adapter name.

2. Take one of the following actions, depending on your configuration:

| If the FC port configuration is... | Then...                                                                                                                                                        |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The same for both ports            | Answer "y" when prompted by the system, because modifying one port in a port pair also modifies the other port.                                                |
| Different                          | a. Answer "n" when prompted by the system.<br>b. Restore the FC port configuration:<br><br><pre>ucadmin modify -m fc -t<br/>initiator targetadapter_name</pre> |

3. Exit Maintenance mode:

```
halt
```

After you issue the command, wait until the system stops at the LOADER prompt.

4. Boot the node back into Maintenance mode for the configuration changes to take effect:

```
boot_ontap maint
```

5. Verify the values of the variables:

```
ucadmin show
```

6. Exit Maintenance mode and display the LOADER prompt:

```
halt
```

### Configuring the FC-to-SAS bridges (MetroCluster FC configurations only)

If you replaced the FC-to-SAS bridges, you must configure them when restoring the MetroCluster configuration. The procedure is identical to the initial configuration of an FC-to-SAS bridge.

#### Steps

1. Power on the FC-to-SAS bridges.
2. Set the IP address on the Ethernet ports by using the `set IPAddress port ipaddress` command.
  - `port` can be either "MP1" or "MP2".
  - `ipaddress` can be an IP address in the format xxx.xxx.xxx.xxx.

In the following example, the IP address is 10.10.10.55 on Ethernet port 1:

```
Ready.
set IPAddress MP1 10.10.10.55

Ready. *
```

3. Set the IP subnet mask on the Ethernet ports by using the `set IPSubnetMask port mask` command.

- `port` can be "MP1" or "MP2".
- `mask` can be a subnet mask in the format `xxx.xxx.xxx.xxx`.

In the following example, the IP subnet mask is 255.255.255.0 on Ethernet port 1:

```
Ready.
set IPSubnetMask MP1 255.255.255.0

Ready. *
```

4. Set the speed on the Ethernet ports by using the `set EthernetSpeed port speed` command.

- `port` can be "MP1" or "MP2".
- `speed` can be "100" or "1000".

In the following example, the Ethernet speed is set to 1000 on Ethernet port 1.

```
Ready.
set EthernetSpeed MP1 1000

Ready. *
```

5. Save the configuration by using the `saveConfiguration` command, and restart the bridge when prompted to do so.

Saving the configuration after configuring the Ethernet ports enables you to proceed with the bridge configuration using Telnet and enables you to access the bridge using FTP to perform firmware updates.

The following example shows the `saveConfiguration` command and the prompt to restart the bridge.

```
Ready.
SaveConfiguration
 Restart is necessary....
 Do you wish to restart (y/n) ?
Confirm with 'y'. The bridge will save and restart with the new
settings.
```

6. After the FC-to-SAS bridge reboots, log in again.

7. Set the speed on the FC ports by using the `set fcdatarate port speed` command.

- port can be "1" or "2".
- speed can be "2 Gb", "4 Gb", "8 Gb", or "16 Gb", depending on your model bridge.

In the following example, the port FC1 speed is set to "8 Gb".

```
Ready.
set fcdatarate 1 8Gb

Ready. *
```

8. Set the topology on the FC ports by using the `set FCConnMode port mode` command.

- port can be "1" or "2".
- mode can be "ptp", "loop", "ptp-loop", or "auto".

In the following example, the port FC1 topology is set to "ptp".

```
Ready.
set FCConnMode 1 ptp

Ready. *
```

9. Save the configuration by using the `saveConfiguration` command, and restart the bridge when prompted to do so.

The following example shows the `saveConfiguration` command and the prompt to restart the bridge.

```
Ready.
SaveConfiguration
 Restart is necessary....
 Do you wish to restart (y/n) ?
Confirm with 'y'. The bridge will save and restart with the new
settings.
```

10. After the FC-to-SAS bridge reboots, log in again.
11. If the FC-to-SAS bridge is running firmware 1.60 or later, enable SNMP.

```
Ready.
set snmp enabled

Ready. *
saveconfiguration

Restart is necessary....
Do you wish to restart (y/n) ?

Verify with 'y' to restart the FibreBridge.
```

12. Power off the FC-to-SAS bridges.

### Configuring the FC switches (MetroCluster FC configurations only)

If you have replaced the FC switches in the disaster site, you must configure them using the vendor-specific procedures. You must configure one switch, verify that storage access on the surviving site is not impacted, and then configure the second switch.

#### Related tasks

[Port assignments for FC switches](#)

#### Configuring a Brocade FC switch after site disaster

You must use this Brocade-specific procedure to configure the replacement switch and enable the ISL ports.

#### About this task

The examples in this procedure are based on the following assumptions:

- Site A is the disaster site.
- FC\_switch\_A\_1 has been replaced.
- FC\_switch\_A\_2 has been replaced.
- Site B is the surviving site.
- FC\_switch\_B\_1 is healthy.
- FC\_switch\_B\_2 is healthy.

You must verify that you are using the specified port assignments when you cable the FC switches:

- [Port assignments for FC switches](#)

The examples show two FC-to-SAS bridges. If you have more bridges, you must disable and subsequently enable the additional ports.

#### Steps

1. Boot and pre-configure the new switch:

- a. Power up the new switch and let it boot up.
- b. Check the firmware version on the switch to confirm it matches the version of the other FC switches:

```
firmwareShow
```

- c. Configure the new switch as described in the following topics, skipping the steps for configuring zoning on the switch.

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

- d. Disable the switch persistently:

```
switchcfgpersistentdisable
```

The switch will remain disabled after a reboot or fastboot. If this command is not available, you should use the `switchdisable` command.

The following example shows the command on BrocadeSwitchA:

```
BrocadeSwitchA:admin> switchcfgpersistentdisable
```

The following example shows the command on BrocadeSwitchB:

```
BrocadeSwitchA:admin> switchcfgpersistentdisable
```

2. Complete configuration of the new switch:

- a. Enable the ISLs on the surviving site:

```
portcfgpersistentenable port-number
```

```
FC_switch_B_1:admin> portcfgpersistentenable 10
FC_switch_B_1:admin> portcfgpersistentenable 11
```

- b. Enable the ISLs on the replacement switches:

```
portcfgpersistentenable port-number
```

```
FC_switch_A_1:admin> portcfgpersistentenable 10
FC_switch_A_1:admin> portcfgpersistentenable 11
```

- c. On the replacement switch (FC\_switch\_A\_1 in this example) verify that the ISL's are online:

```
switchshow
```

```
FC_switch_A_1:admin> switchshow
switchName: FC_switch_A_1
switchType: 71.2
switchState:Online
switchMode: Native
switchRole: Principal
switchDomain: 4
switchId: fffc03
switchWwn: 10:00:00:05:33:8c:2e:9a
zoning: OFF
switchBeacon: OFF

Index Port Address Media Speed State Proto
=====
...
10 10 030A00 id 16G Online FC E-Port
10:00:00:05:33:86:89:cb "FC_switch_A_1"
11 11 030B00 id 16G Online FC E-Port
10:00:00:05:33:86:89:cb "FC_switch_A_1" (downstream)
...
```

3. Persistently enable the switch:

```
switchcfgpersistentenable
```

4. Verify that the ports are online:

```
switchshow
```

### Configuring a Cisco FC switch after site disaster

You must use the Cisco-specific procedure to configure the replacement switch and enable the ISL ports.

#### About this task

The examples in this procedure are based on the following assumptions:

- Site A is the disaster site.
- FC\_switch\_A\_1 has been replaced.
- FC\_switch\_A\_2 has been replaced.
- Site B is the surviving site.
- FC\_switch\_B\_1 is healthy.
- FC\_switch\_B\_2 is healthy.

#### Steps

1. Configure the switch:

- a. Refer to [Fabric-attached MetroCluster installation and configuration](#)
- b. Follow the steps for configuring the switch in [Configuring the Cisco FC switches](#) section, *except* for the "Configuring zoning on a Cisco FC switch" section:

Zoning is configured later in this procedure.

2. On the healthy switch (in this example, FC\_switch\_B\_1), enable the ISL ports.

The following example shows the commands to enable the ports:

```
FC_switch_B_1# conf t
FC_switch_B_1(config)# int fc1/14-15
FC_switch_B_1(config)# no shut
FC_switch_B_1(config)# end
FC_switch_B_1# copy running-config startup-config
FC_switch_B_1#
```

3. Verify that the ISL ports are up by using the show interface brief command.

4. Retrieve the zoning information from the fabric.

The following example shows the commands to distribute the zoning configuration:

```
FC_switch_B_1(config-zone)# zoneset distribute full vsan 10
FC_switch_B_1(config-zone)# zoneset distribute full vsan 20
FC_switch_B_1(config-zone)# end
```

FC\_switch\_B\_1 is distributed to all other switches in the fabric for "vsan 10" and "vsan 20", and the zoning information is retrieved from FC\_switch\_A\_1.

5. On the healthy switch, verify that the zoning information is properly retrieved from the partner switch:

```
show zone
```

```

FC_switch_B_1# show zone
zone name FC-VI_Zone_1_10 vsan 10
 interface fc1/1 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/1 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25A vsan 20
 interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25B vsan 20
 interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
 interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0
FC_switch_B_1#

```

6. Determine the worldwide names (WWNs) of the switches in the switch fabric.

In this example, the two switch WWNs are as follows:

- FC\_switch\_A\_1: 20:00:54:7f:ee:b8:24:c0
- FC\_switch\_B\_1: 20:00:54:7f:ee:c6:80:78

```

FC_switch_B_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:c6:80:78
FC_switch_B_1#

```

```

FC_switch_A_1# show wwn switch
Switch WWN is 20:00:54:7f:ee:b8:24:c0
FC_switch_A_1#

```

7. Enter configuration mode for the zone and remove zone members that do not belong to the switch WWNs of the two switches:

```
no member interface interface-ide swwn wwn
```

In this example, the following members are not associated with the WWN of either of the switches in the fabric and must be removed:

- Zone name FC-VI\_Zone\_1\_10 vsan 10
  - Interface fc1/1 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/2 swwn 20:00:54:7f:ee:e3:86:50



AFF A700 and FAS9000 systems support four FC-VI ports. You must remove all four ports from the FC-VI zone.

- Zone name STOR\_Zone\_1\_20\_25A vsan 20
  - Interface fc1/5 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50
- Zone name STOR\_Zone\_1\_20\_25B vsan 20
  - Interface fc1/8 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/9 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/10 swwn 20:00:54:7f:ee:e3:86:50
  - Interface fc1/11 swwn 20:00:54:7f:ee:e3:86:50

The following example shows the removal of these interfaces:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# zone name FC-VI_Zone_1_10 vsan 10
FC_switch_B_1(config-zone)# no member interface fc1/1 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/2 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25A vsan 20
FC_switch_B_1(config-zone)# no member interface fc1/5 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/8 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/9 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/10 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/11 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25B vsan 20
FC_switch_B_1(config-zone)# no member interface fc1/8 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/9 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/10 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# no member interface fc1/11 swwn
20:00:54:7f:ee:e3:86:50
FC_switch_B_1(config-zone)# save running-config startup-config
FC_switch_B_1(config-zone)# zoneset distribute full 10
FC_switch_B_1(config-zone)# zoneset distribute full 20
FC_switch_B_1(config-zone)# end
FC_switch_B_1# copy running-config startup-config

```

## 8. Add the ports of the new switch to the zones.

The following example assumes that the cabling on the replacement switch is the same as on the old switch:

```

FC_switch_B_1# conf t
FC_switch_B_1(config)# zone name FC-VI_Zone_1_10 vsan 10
FC_switch_B_1(config-zone)# member interface fc1/1 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/2 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25A vsan 20
FC_switch_B_1(config-zone)# member interface fc1/5 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/8 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/9 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/10 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/11 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# zone name STOR_Zone_1_20_25B vsan 20
FC_switch_B_1(config-zone)# member interface fc1/8 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/9 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/10 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# member interface fc1/11 swwn
20:00:54:7f:ee:c6:80:78
FC_switch_B_1(config-zone)# save running-config startup-config
FC_switch_B_1(config-zone)# zoneset distribute full 10
FC_switch_B_1(config-zone)# zoneset distribute full 20
FC_switch_B_1(config-zone)# end
FC_switch_B_1# copy running-config startup-config

```

9. Verify that the zoning is properly configured: show zone

The following example output shows the three zones:

```

FC_switch_B_1# show zone
zone name FC-VI_Zone_1_10 vsan 10
 interface fc1/1 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/2 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/1 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/2 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25A vsan 20
 interface fc1/5 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/8 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/9 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/10 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/11 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0

zone name STOR_Zone_1_20_25B vsan 20
 interface fc1/8 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/9 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/10 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/11 swwn 20:00:54:7f:ee:c6:80:78
 interface fc1/5 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/8 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/9 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/10 swwn 20:00:54:7f:ee:b8:24:c0
 interface fc1/11 swwn 20:00:54:7f:ee:b8:24:c0
FC_switch_B_1#

```

## Verifying the storage configuration

You must confirm that all storage is visible from the surviving nodes.

### Steps

1. Confirm that all storage components at the disaster site are the same in quantity and type at the surviving site.

The surviving site and disaster site should have the same number of disk shelf stacks, disk shelves, and disks. In a bridge-attached or fabric-attached MetroCluster configuration, the sites should have the same number of FC-to-SAS bridges.

2. Confirm that all disks that have been replaced at the disaster site are unowned:

```
run local disk show-n
```

Disks should appear as being unowned.

3. If no disks were replaced, confirm that all disks are present:

```
disk show
```

### Powering on the equipment at the disaster site

You must power on the MetroCluster components at the disaster site when you are ready to prepare for switchback. In addition, you must also recable the SAS storage connections in direct-attached MetroCluster configurations and enable non-Inter-Switch Link ports in fabric-attached MetroCluster configurations.

#### Before you begin

You must have already replaced and cabled the MetroCluster components exactly as the old ones.

#### [Fabric-attached MetroCluster installation and configuration](#)

#### [Stretch MetroCluster installation and configuration](#)

#### About this task

The examples in this procedure assume the following:

- Site A is the disaster site.
  - FC\_switch\_A\_1 has been replaced.
  - FC\_switch\_A\_2 has been replaced.
- Site B is the surviving site.
  - FC\_switch\_B\_1 is healthy.
  - FC\_switch\_B\_2 is healthy.

The FC switches are present only in fabric-attached MetroCluster configurations.

#### Steps

1. In a stretch MetroCluster configuration using SAS cabling (and no FC switch fabric or FC-to-SAS bridges), connect all the storage including the remote storage across both sites.

The controller at the disaster site must remain powered off or at the LOADER prompt.

2. On the surviving site, disable disk autoassignment:

```
storage disk option modify -autoassign off *
```

```
cluster_B::> storage disk option modify -autoassign off *
2 entries were modified.
```

3. On the surviving site, confirm that disk autoassignment is off:

```
storage disk option show
```

```

cluster_B::> storage disk option show
Node BKg. FW. Upd. Auto Copy Auto Assign Auto Assign Policy

node_B_1 on on off default
node_B_2 on on off default
2 entries were displayed.

cluster_B::>

```

4. Turn on the disk shelves at the disaster site and make sure that all disks are running.
5. In a bridge-attached or fabric-attached MetroCluster configuration, turn on all FC-to-SAS bridges at the disaster site.
6. If any disks were replaced, leave the controllers powered off or at the LOADER prompt.
7. In a fabric-attached MetroCluster configuration, enable the non-ISL ports on the FC switches.

|                                   |                                                    |
|-----------------------------------|----------------------------------------------------|
| <b>If the switch vendor is...</b> | <b>Then use these steps to enable the ports...</b> |
|-----------------------------------|----------------------------------------------------|

## Brocade

- a. Persistently enable the ports connected to the FC-to-SAS bridges: `portpersistentenable port-number`

In the following example, ports 6 and 7 are enabled:

```
FC_switch_A_1:admin>
portpersistentenable 6
FC_switch_A_1:admin>
portpersistentenable 7

FC_switch_A_1:admin>
```

- b. Persistently enable the ports connected to the HBAs and FC-VI adapters:

`portpersistentenable port-number`

In the following example, ports 6 and 7 are enabled:

```
FC_switch_A_1:admin>
portpersistentenable 1
FC_switch_A_1:admin>
portpersistentenable 2
FC_switch_A_1:admin>
portpersistentenable 4
FC_switch_A_1:admin>
portpersistentenable 5
FC_switch_A_1:admin>
```



For AFF A700 and FAS9000 systems, you must persistently enable all four FC-VI ports by using the `switchcfgpersistentenable` command.

- c. Repeat substeps a and b for the second FC switch at the surviving site.

Cisco

- a. Enter configuration mode for the interface, and then enable the ports with the no shut command.

In the following example, port fc1/36 is disabled:

```
FC_switch_A_1# conf t
FC_switch_A_1(config)#
interface fc1/36
FC_switch_A_1(config)# no shut
FC_switch_A_1(config-if)# end
FC_switch_A_1# copy running-
config startup-config
```

- b. Verify that the switch port is enabled: show interface brief
- c. Repeat Substeps a and b on the other ports connected to the FC-to-SAS bridges, HBAs, and FC-VI adapters.
- d. Repeat Substeps a, b, and c for the second FC switch at the surviving site.

## Assigning ownership for replaced drives

If you replaced drives when restoring hardware at the disaster site or you had to zero drives or remove ownership, you must assign ownership to the affected drives.

### Before you begin

The disaster site must have at least as many available drives as it did prior to the disaster.

The drives shelves and drives arrangement must meet the requirements in [Required MetroCluster IP component and naming conventions](#) section of the [MetroCluster IP installation and configuration](#).

### About this task

These steps are performed on the cluster at the disaster site.

This procedure shows the reassignment of all drives and the creation of new plexes at the disaster site. The new plexes are remote plexes of surviving site and local plexes of disaster site.

This section provides examples for two and four-node configurations. For two-node configurations, you can ignore references to the second node at each site. For eight-node configurations, you must account for the additional nodes on the second DR group. The examples make the following assumptions:

- Site A is the disaster site.
  - node\_A\_1 has been replaced.
  - node\_A\_2 has been replaced.

Present only in four-node MetroCluster configurations.

- Site B is the surviving site.
  - node\_B\_1 is healthy.
  - node\_B\_2 is healthy.

Present only in four-node MetroCluster configurations.

The controller modules have the following original system IDs:

| Number of nodes in MetroCluster configuration | Node       | Original system ID |
|-----------------------------------------------|------------|--------------------|
| Four                                          | node_A_1   | 4068741258         |
| node_A_2                                      | 4068741260 | node_B_1           |
| 4068741254                                    | node_B_2   | 4068741256         |
| Two                                           | node_A_1   | 4068741258         |

You should keep in mind the following points when assigning the drives:

- The old-count-of-disks must be at least the same number of disks for each node that were present before the disaster.

If a lower number of disks is specified or present, the healing operations might not be completed due to insufficient space.

- The new plexes to be created are remote plexes belonging to the surviving site (node\_B\_x pool1) and local plexes belonging to the disaster site (node\_B\_x pool0).
- The total number of required drives should not include the root aggr disks.

If n disks are assigned to pool1 of the surviving site, then n-3 disks should be assigned to the disaster site with the assumption that the root aggregate uses three disks.

- None of the disks can be assigned to a pool that is different from the one to which all other disks on the same stack are assigned.
- Disks belonging to the surviving site are assigned to pool 1 and disks belonging to the disaster site are assigned to pool 0.

## Steps

1. Assign the new, unowned drives based on whether you have a four-node or two-node MetroCluster configuration:
  - For four-node MetroCluster configurations, assign the new, unowned disks to the appropriate disk pools by using the following series of commands on the replacement nodes:
    - i. Systematically assign the replaced disks for each node to their respective disk pools:

```
disk assign -s sysid -n old-count-of-disks -p pool
```

From the surviving site, you issue a disk assign command for each node:

```
cluster_B::> disk assign -s node_B_1-sysid -n old-count-of-disks
-p 1 **\ (remote pool of surviving site\)**
cluster_B::> disk assign -s node_B_2-sysid -n old-count-of-disks
-p 1 **\ (remote pool of surviving site\)**
cluster_B::> disk assign -s node_A_1-old-sysid -n old-count-of-
disks -p 0 **\ (local pool of disaster site\)**
cluster_B::> disk assign -s node_A_2-old-sysid -n old-count-of-
disks -p 0 **\ (local pool of disaster site\)**
```

The following example shows the commands with the system IDs:

```
cluster_B::> disk assign -s 4068741254 -n 21 -p 1
cluster_B::> disk assign -s 4068741256 -n 21 -p 1
cluster_B::> disk assign -s 4068741258 -n 21 -p 0
cluster_B::> disk assign -s 4068741260 -n 21 -p 0
```

ii. Confirm the ownership of the disks:

```
storage disk show -fields owner, pool
```

```

storage disk show -fields owner, pool
cluster_A::> storage disk show -fields owner, pool
disk owner pool
----- -
0c.00.1 node_A_1 Pool0
0c.00.2 node_A_1 Pool0
.
.
.
0c.00.8 node_A_1 Pool1
0c.00.9 node_A_1 Pool1
.
.
.
0c.00.15 node_A_2 Pool0
0c.00.16 node_A_2 Pool0
.
.
.
0c.00.22 node_A_2 Pool1
0c.00.23 node_A_2 Pool1
.
.
.

```

- For two-node MetroCluster configurations, assign the new, unowned disks to the appropriate disk pools by using the following series of commands on the replacement node:

- i. Display the local shelf IDs:

```
run local storage show shelf
```

- ii. Assign the replaced disks for the healthy node to pool 1:

```
run local disk assign -shelf shelf-id -n old-count-of-disks -p 1 -s
node_B_1-sysid -f
```

- iii. Assign the replaced disks for the replacement node to pool 0:

```
run local disk assign -shelf shelf-id -n old-count-of-disks -p 0 -s
node_A_1-sysid -f
```

2. On the surviving site, turn on automatic disk assignment again:

```
storage disk option modify -autoassign on *
```

```
cluster_B::> storage disk option modify -autoassign on *
2 entries were modified.
```

3. On the surviving site, confirm that automatic disk assignment is on:

```
storage disk option show
```

```
cluster_B::> storage disk option show
Node BKg. FW. Upd. Auto Copy Auto Assign Auto Assign Policy
----- -
node_B_1 on on on default
node_B_2 on on on default
2 entries were displayed.

cluster_B::>
```

### Related information

[Disk and aggregate management](#)

[How MetroCluster configurations use SyncMirror to provide data redundancy](#)

### Performing aggregate healing and restoring mirrors (MetroCluster FC configurations)

After replacing hardware and assigning disks, you can perform the MetroCluster healing operations. You must then confirm that aggregates are mirrored and, if necessary, restart mirroring.

#### Steps

1. Perform the two phases of healing (aggregate healing and root healing) on the disaster site:

```
cluster_B::> metrocluster heal -phase aggregates

cluster_B::> metrocluster heal -phase root-aggregates
```

2. Monitor the healing and verify that the aggregates are in either the resyncing or mirrored state:

```
storage aggregate show -node local
```

| If the aggregate shows this state... | Then...                                                      |
|--------------------------------------|--------------------------------------------------------------|
| resyncing                            | No action is required. Let the aggregate complete resyncing. |

|                  |                                                                                                                       |
|------------------|-----------------------------------------------------------------------------------------------------------------------|
| mirror degraded  | Proceed to <a href="#">If one or more plexes remain offline, additional steps are required to rebuild the mirror.</a> |
| mirrored, normal | No action is required.                                                                                                |
| unknown, offline | The root aggregate shows this state if all the disks on the disaster sites were replaced.                             |

```

cluster_B::> storage aggregate show -node local

Aggregate Size Available Used% State #Vols Nodes RAID
Status

node_B_1_aggr1
 227.1GB 11.00GB 95% online 1 node_B_1 raid_dp,
resyncing
NodeA_1_aggr2
 430.3GB 28.02GB 93% online 2 node_B_1 raid_dp,
mirror
degraded
node_B_1_aggr3
 812.8GB 85.37GB 89% online 5 node_B_1 raid_dp,
mirrored,
normal

3 entries were displayed.

cluster_B::>

```

In the following examples, the three aggregates are each in a different state:

| Node           | State            |
|----------------|------------------|
| node_B_1_aggr1 | resyncing        |
| node_B_1_aggr2 | mirror degraded  |
| node_B_1_aggr3 | mirrored, normal |

- If one or more plexes remain offline, additional steps are required to rebuild the mirror.

In the preceding table, the mirror for node\_B\_1\_aggr2 must be rebuilt.

- View details of the aggregate to identify any failed plexes:

```
storage aggregate show -r -aggregate node_B_1_aggr2
```

In the following example, plex /node\_B\_1\_aggr2/plex0 is in a failed state:

```
cluster_B::> storage aggregate show -r -aggregate node_B_1_aggr2

Owner Node: node_B_1
Aggregate: node_B_1_aggr2 (online, raid_dp, mirror degraded) (block
checksums)
Plex: /node_B_1_aggr2/plex0 (offline, failed, inactive, pool0)
RAID Group /node_B_1_aggr2/plex0/rg0 (partial)
Usable
Physical
Position Disk Pool Type RPM Size
Size Status

Plex: /node_B_1_aggr2/plex1 (online, normal, active, pool1)
RAID Group /node_B_1_aggr2/plex1/rg0 (normal, block checksums)
Usable
Physical
Position Disk Pool Type RPM Size
Size Status

dparity 1.44.8 1 SAS 15000 265.6GB
273.5GB (normal)
parity 1.41.11 1 SAS 15000 265.6GB
273.5GB (normal)
data 1.42.8 1 SAS 15000 265.6GB
273.5GB (normal)
data 1.43.11 1 SAS 15000 265.6GB
273.5GB (normal)
data 1.44.9 1 SAS 15000 265.6GB
273.5GB (normal)
data 1.43.18 1 SAS 15000 265.6GB
273.5GB (normal)
6 entries were displayed.

cluster_B::>
```

b. Delete the failed plex:

```
storage aggregate plex delete -aggregate aggregate-name -plex plex
```

c. Reestablish the mirror:

```
storage aggregate mirror -aggregate aggregate-name
```

- d. Monitor the resynchronization and mirroring status of the plex until all mirrors are reestablished and all aggregates show mirrored, normal status:

```
storage aggregate show
```

### Reassigning disk ownership for root aggregates to replacement controller modules (MetroCluster FC configurations)

If one or both of the controller modules or NVRAM cards were replaced at the disaster site, the system ID has changed and you must reassign disks belonging to the root aggregates to the replacement controller modules.

#### About this task

Because the nodes are in switchover mode and healing has been done, only the disks containing the root aggregates of pool1 of the disaster site will be reassigned in this section. They are the only disks still owned by the old system ID at this point.

This section provides examples for two and four-node configurations. For two-node configurations, you can ignore references to the second node at each site. For eight-node configurations, you must account for the additional nodes on the second DR group. The examples make the following assumptions:

- Site A is the disaster site.
  - node\_A\_1 has been replaced.
  - node\_A\_2 has been replaced.

Present only in four-node MetroCluster configurations.

- Site B is the surviving site.
  - node\_B\_1 is healthy.
  - node\_B\_2 is healthy.

Present only in four-node MetroCluster configurations.

The old and new system IDs were identified in [Replace hardware and boot new controllers](#).

The examples in this procedure use controllers with the following system IDs:

| Number of nodes | Node     | Original system ID | New system ID |
|-----------------|----------|--------------------|---------------|
| Four            | node_A_1 | 4068741258         | 1574774970    |
|                 | node_A_2 | 4068741260         | 1574774991    |
|                 | node_B_1 | 4068741254         | unchanged     |
|                 | node_B_2 | 4068741256         | unchanged     |

|     |          |            |            |
|-----|----------|------------|------------|
| Two | node_A_1 | 4068741258 | 1574774970 |
|-----|----------|------------|------------|

## Steps

1. With the replacement node in Maintenance mode, reassign the root aggregate disks:

```
disk reassign -s old-system-ID -d new-system-ID
```

```
*> disk reassign -s 4068741258 -d 1574774970
```

2. View the disks to confirm the ownership change of the pool1 root aggr disks of the disaster site to the replacement node:

```
disk show
```

The output might show more or fewer disks, depending on how many disks are in the root aggregate and whether any of these disks failed and were replaced. If the disks were replaced, then Pool0 disks will not appear in the output.

The pool1 root aggregate disks of the disaster site should now be assigned to the replacement node.

```

*> disk show
Local System ID: 1574774970

 DISK OWNER POOL SERIAL NUMBER HOME
DR HOME

sw_A_1:6.126L19 node_A_1(1574774970) Pool0 serial-number
node_A_1(1574774970)
sw_A_1:6.126L3 node_A_1(1574774970) Pool0 serial-number
node_A_1(1574774970)
sw_A_1:6.126L7 node_A_1(1574774970) Pool0 serial-number
node_A_1(1574774970)
sw_B_1:6.126L8 node_A_1(1574774970) Pool11 serial-number
node_A_1(1574774970)
sw_B_1:6.126L24 node_A_1(1574774970) Pool11 serial-number
node_A_1(1574774970)
sw_B_1:6.126L2 node_A_1(1574774970) Pool11 serial-number
node_A_1(1574774970)

*> aggr status
 Aggr State Status
node_A_1_root online raid_dp, aggr
 mirror degraded
 64-bit

*>

```

### 3. View the aggregate status:

```
aggr status
```

The output might show more or fewer disks, depending on how many disks are in the root aggregate and whether any of these disks failed and were replaced. If disks were replaced, then Pool0 disks will not appear in output.

```

*> aggr status
 Aggr State Status
node_A_1_root online raid_dp, aggr
 mirror degraded
 64-bit

*>

```

### 4. Delete the contents of the mailbox disks:

```
mailbox destroy local
```

5. If the aggregate is not online, bring it online:

```
aggr online aggr_name
```

6. Halt the node to display the LOADER prompt:

```
halt
```

## Booting the new controller modules (MetroCluster FC configurations)

After aggregate healing has been completed for both the data and root aggregates, you must boot the node or nodes at the disaster site.

### About this task

This task begins with the nodes showing the LOADER prompt.

### Steps

1. Display the boot menu:

```
boot_ontap menu
```

2. From the boot menu, select option 6, **Update flash from backup config**.
3. Respond `y` to the following prompt:

```
This will replace all flash-based configuration with the last backup to disks.
Are you sure you want to continue?: y
```

The system will boot twice, the second time to load the new configuration.



If you did not clear the NVRAM contents of a used replacement controller, then you might see a panic with the following message:

```
PANIC: NVRAM contents are invalid...
```

If this occurs, repeat [From the boot menu, select option 6, Update flash from backup config](#). to boot the system to the ONTAP prompt. You then need to [Reset the boot recovery and rdb\\_corrupt bootargs](#)

4. Mirror the root aggregate on plex 0:
  - a. Assign three pool0 disks to the new controller module.
  - b. Mirror the root aggregate pool1 plex:

```
aggr mirror root-aggr-name
```

- c. Assign unowned disks to pool0 on the local node
5. If you have a four-node configuration, repeat the previous steps on the other node at the disaster site.
  6. Refresh the MetroCluster configuration:
    - a. Enter advanced privilege mode:

```
set -privilege advanced
```

b. Refresh the configuration:

```
metrocluster configure -refresh true
```

c. Return to admin privilege mode:

```
set -privilege admin
```

7. Confirm that the replacement nodes at the disaster site are ready for switchback:

```
metrocluster node show
```

The replacement nodes should be in “waiting for switchback recovery” mode. If they are in “normal” mode instead, you can reboot the replacement nodes. After that boot, the nodes should be in “waiting for switchback recovery” mode.

The following example shows that the replacement nodes are ready for switchback:

```
cluster_B::> metrocluster node show
DR Configuration DR
Grp Cluster Node State Mirroring Mode

1 cluster_B
 node_B_1 configured enabled switchover completed
 node_B_2 configured enabled switchover completed
 cluster_A
 node_A_1 configured enabled waiting for switchback
recovery
 node_A_2 configured enabled waiting for switchback
recovery
4 entries were displayed.

cluster_B::>
```

### What to do next

Proceed to [Complete the disaster recovery process](#).

### Reset the boot\_recovery and rdb\_corrupt bootargs

If required, you can reset the boot\_recovery and rdb\_corrupt\_bootargs

### Steps

1. Halt the node back to the LOADER prompt:

```
siteA::*> halt -node <node-name>
```

2. Check if the following bootargs have been set:

```
LOADER> printenv bootarg.init.boot_recovery
LOADER> printenv bootarg.rdb_corrupt
```

3. If either bootarg has been set to a value, unset it and boot ONTAP:

```
LOADER> unsetenv bootarg.init.boot_recovery
LOADER> unsetenv bootarg.rdb_corrupt
LOADER> saveenv
LOADER> bye
```

## Preparing for switchback in a mixed configuration (recovery during transition)

You must perform certain tasks in order to prepare the mixed MetroCluster IP and FC configuration for the switchback operation. This procedure only applies to configurations that encountered a failure during the MetroCluster FC to IP transition process.

### About this task

This procedure should only be used when performing recovery on a system that was in mid-transition when the failure occurred.

In this scenario, the MetroCluster is a mixed configuration:

- One DR group consists of fabric-attached MetroCluster FC nodes.  
You must perform the MetroCluster FC recovery steps on these nodes.
- One DR group consists of MetroCluster IP nodes.  
You must perform the MetroCluster IP recovery steps on these nodes.

### Steps

Perform the steps in the following order.

1. Prepare the FC nodes for switchback by performing the following tasks in order:
  - a. [Verifying port configuration \(MetroCluster FC configurations only\)](#)
  - b. [Configuring the FC-to-SAS bridges \(MetroCluster FC configurations only\)](#)
  - c. [Configuring the FC switches \(MetroCluster FC configurations only\)](#)
  - d. [Verifying the storage configuration](#) (only perform these steps on replaced drives on the MetroCluster FC nodes)
  - e. [Powering on the equipment at the disaster site](#) (only perform these steps on replaced drives on the MetroCluster FC nodes)
  - f. [Assigning ownership for replaced drives](#) (only perform these steps on replaced drives on the MetroCluster FC nodes)
  - g. Perform the steps in [Reassigning disk ownership for root aggregates to replacement controller modules \(MetroCluster FC configurations\)](#), up to and including the step to issue the mailbox destroy command.

h. Destroy the local plex (plex 0) of the root aggregate:

```
aggr destroy plex-id
```

i. If the root aggr is not online, bring it online.

2. Boot the MetroCluster FC nodes.

You must perform these steps on both of the MetroCluster FC nodes.

a. Display the boot menu:

```
boot_ontap menu
```

b. From the boot menu, select option 6, **Update flash from backup config**.

c. Respond *y* to the following prompt:

```
This will replace all flash-based configuration with the last backup to
disks. Are you sure you want to continue?: y
```

The system will boot twice, the second time to load the new configuration.



If you did not clear the NVRAM contents of a used replacement controller, then you might see a panic with the following message: `PANIC: NVRAM contents are invalid...` If this occurs, repeat these substeps to boot the system to the ONTAP prompt. You then need to [Reset the boot recovery and rdb\\_corrupt bootargs](#)

3. Mirror the root aggregate on plex 0:

You must perform these steps on both of the MetroCluster FC nodes.

a. Assign three pool0 disks to the new controller module.

b. Mirror the root aggregate pool1 plex:

```
aggr mirror root-aggr-name
```

c. Assign unowned disks to pool0 on the local node

4. Return to Maintenance mode.

You must perform these steps on both of the MetroCluster FC nodes.

a. Halt the node:

```
halt
```

b. Boot the node to Maintenance mode:

```
boot_ontap maint
```

5. Delete the contents of the mailbox disks:

```
mailbox destroy local
```

You must perform these steps on both of the MetroCluster FC nodes.

6. Halt the nodes:

```
halt
```

7. After the nodes boot up, verify the status of the node:

```
metrocluster node show
```

```
siteA::*> metrocluster node show
DR
Group Cluster Node Configuration DR
State Mirroring Mode

1 siteA
 wmc66-a1 configured enabled waiting for
switchback recovery
 wmc66-a2 configured enabled waiting for
switchback recovery
 siteB
 wmc66-b1 configured enabled switchover
completed
 wmc66-b2 configured enabled switchover
completed
2 siteA
 wmc55-a1 - - -
 wmc55-a2 unreachable - -
 siteB
 wmc55-b1 configured enabled switchover
completed
 wmc55-b2 configured
```

8. Prepare the MetroCluster IP nodes for switchback by performing the tasks in [Preparing for switchback in a MetroCluster IP configuration](#) up to and including [Deleting failed plexes owned by the surviving site \(MetroCluster IP configurations\)](#).
9. On the MetroCluster FC nodes, perform the steps in [Performing aggregate healing and restoring mirrors \(MetroCluster FC configurations\)](#).
10. On the MetroCluster IP nodes, perform the steps in [Performing aggregate healing and restoring mirrors \(MetroCluster IP configurations\)](#).
11. Proceed through the remaining tasks of the recovery process beginning with [Reestablishing object stores for FabricPool configurations](#).

**Reset the boot\_recovery and rdb\_corrupt bootargs**

If required, you can reset the boot\_recovery and rdb\_corrupt\_bootargs

**Steps**

1. Halt the node back to the LOADER prompt:

```
siteA::*> halt -node <node-name>
```

2. Check if the following bootargs have been set:

```
LOADER> printenv bootarg.init.boot_recovery
LOADER> printenv bootarg.rdb_corrupt
```

3. If either bootarg has been set to a value, unset it and boot ONTAP:

```
LOADER> unsetenv bootarg.init.boot_recovery
LOADER> unsetenv bootarg.rdb_corrupt
LOADER> saveenv
LOADER> bye
```

## Completing recovery

Perform the required tasks to complete the recovery from a multi-controller or storage failure.

### Reestablishing object stores for FabricPool configurations

If one of the object stores in a FabricPool mirror was co-located with the MetroCluster disaster site and was destroyed, you must reestablish the object store and the FabricPool mirror.

#### About this task

- If the object-stores are remote and a MetroCluster site is destroyed, you do not need to rebuild the object store, and the original object store configurations as well as cold data contents are retained.
- For more information about FabricPool configurations, see the [Disk and aggregates management](#).

#### Step

1. Follow the procedure "Replacing a FabricPool mirror on a MetroCluster configuration" in the [Disk and aggregates management](#).

### Verifying licenses on the replaced nodes

You must install new licenses for the replacement nodes if the impaired nodes were using ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

#### About this task

Until you install license keys, features requiring standard licenses continue to be available to the replacement node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on

the replacement node as soon as possible.

The licenses keys must be in the 28-character format.

You have a 90-day grace period in which to install the license keys. After the grace period, all old licenses are invalidated. After a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.



If all nodes at a site have been replaced (a single node in the case of a two-node MetroCluster configuration), license keys must be installed on the replacement node or nodes prior to switchback.

## Steps

1. Identify the licenses on the node:

```
license show
```

The following example displays the information about licenses in the system:

```
cluster_B::> license show
 (system license show)

Serial Number: 1-80-00050
Owner: sitel-01
Package Type Description Expiration
----- -
Base license Cluster Base License -
NFS site NFS License -
CIFS site CIFS License -
iSCSI site iSCSI License -
FCP site FCP License -
FlexClone site FlexClone License -

6 entries were displayed.
```

2. Verify that the licenses are good for the node after switchback:

```
metrocluster check license show
```

The following example displays the licenses that are good for the node:

```
cluster_B::> metrocluster check license show
```

| Cluster   | Check                       | Result         |
|-----------|-----------------------------|----------------|
| -----     | -----                       | -----          |
| Cluster_B | negotiated-switchover-ready | not-applicable |
| NFS       | switchback-ready            | not-applicable |
| CIFS      | job-schedules               | ok             |
| iSCSI     | licenses                    | ok             |
| FCP       | periodic-check-enabled      | ok             |

3. If you need new license keys, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.



The new license keys that you require are automatically generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, refer to the *"Who to contact if I have issues with my Licenses?"* section in the Knowledge Base article [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#).

4. Install each license key:

```
system license add -license-code license-key, license-key...+
```

5. Remove the old licenses, if desired:

- a. Check for unused licenses:

```
license clean-up -unused -simulate
```

- b. If the list looks correct, remove the unused licenses:

```
license clean-up -unused
```

## Restoring key management

If data volumes are encrypted, you must restore key management. If the root volume is encrypted, you must recover key management.

### Steps

1. If data volumes are encrypted, restore the keys using the correct command for your key management configuration.

| If you are using...           | Use this command...                                                                                                                             |
|-------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Onboard key management</b> | <pre>security key-manager onboard sync</pre> <p>For more information, see <a href="#">Restoring onboard key management encryption keys</a>.</p> |

|                                |                                                                                                                                                               |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>External key management</b> | <pre>security key-manager key query -node node-name</pre> <p>For more information, see <a href="#">Restoring external key management encryption keys</a>.</p> |
|--------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|

- If the root volume is encrypted, use the procedure in [Recovering key management if the root volume is encrypted](#).

## Performing a switchback

After you heal the MetroCluster configuration, you can perform the MetroCluster switchback operation. The MetroCluster switchback operation returns the configuration to its normal operating state, with the sync-source storage virtual machines (SVMs) on the disaster site active and serving data from the local disk pools.

### Before you begin

- The disaster cluster must have successfully switched over to the surviving cluster.
- Healing must have been performed on the data and root aggregates.
- The surviving cluster nodes must not be in the HA failover state (all nodes must be up and running for each HA pair).
- The disaster site controller modules must be completely booted and not in the HA takeover mode.
- The root aggregate must be mirrored.
- The Inter-Switch Links (ISLs) must be online.
- Any required licenses must be installed on the system.

### Steps

- Confirm that all nodes are in the enabled state:

```
metrocluster node show
```

The following example displays the nodes that are in the enabled state:

```
cluster_B::> metrocluster node show

DR
Group Cluster Node Configuration DR

1 cluster_A
 node_A_1 configured enabled heal roots completed
 node_A_2 configured enabled heal roots completed
 cluster_B
 node_B_1 configured enabled waiting for
switchback recovery
 node_B_2 configured enabled waiting for
switchback recovery
4 entries were displayed.
```

2. Confirm that resynchronization is complete on all SVMs:

```
metrocluster vserver show
```

3. Verify that any automatic LIF migrations being performed by the healing operations have been successfully completed:

```
metrocluster check lif show
```

4. Perform the switchback by running the `metrocluster switchback` command from any node in the surviving cluster.

5. Check the progress of the switchback operation:

```
metrocluster show
```

The switchback operation is still in progress when the output displays "waiting-for-switchback":

```
cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
 Mode switchover
 AUSO Failure Domain -
Remote: cluster_A Configuration state configured
 Mode waiting-for-switchback
 AUSO Failure Domain -
```

The switchback operation is complete when the output displays "normal":

```
cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
 Mode normal
 AUSO Failure Domain -
Remote: cluster_A Configuration state configured
 Mode normal
 AUSO Failure Domain -
```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the following command at the advanced privilege level:

```
metrocluster config-replication resync-status show
```

6. Reestablish any SnapMirror or SnapVault configurations.

In ONTAP 8.3, you need to manually reestablish a lost SnapMirror configuration after a MetroCluster switchback operation. In ONTAP 9.0 and later, the relationship is reestablished automatically.

## Verifying a successful switchback

After performing the switchback, you want to confirm that all aggregates and storage virtual machines (SVMs) are switched back and online.

### Steps

1. Verify that the switched-over data aggregates are switched back:

```
storage aggregate show
```

In the following example, `aggr_b2` on node B2 has switched back:

```
node_B_1::> storage aggregate show
Aggregate Size Available Used% State #Vols Nodes RAID
Status

...
aggr_b2 227.1GB 227.1GB 0% online 0 node_B_2 raid_dp,
mirrored,
normal

node_A_1::> aggr show
Aggregate Size Available Used% State #Vols Nodes RAID
Status

...
aggr_b2 - - - unknown - node_A_1
```

If the disaster site included unmirrored aggregates and the unmirrored aggregates are no longer present, the aggregate might show up with a state of “unknown” in the output of the `storage aggregate show` command. Contact technical support to remove the out-of-date entries for the unmirrored aggregates, reference the Knowledge Base article [How to remove stale unmirrored aggregate entries in a MetroCluster following disaster where storage was lost](#).

2. Verify that all sync-destination SVMs on the surviving cluster are dormant (showing an operational state of “stopped”):

```
vserver show -subtype sync-destination
```

```

node_B_1::> vserver show -subtype sync-destination
 Admin Operational Root
Vserver Type Subtype State State Volume
Aggregate

...
cluster_A-vs1a-mc data sync-destination
 running stopped vs1a_vol aggr_b2

```

Sync-destination aggregates in the MetroCluster configuration have the suffix “-mc” automatically appended to their name to help identify them.

3. Verify the sync-source SVMs on the disaster cluster are up and running:

```
vserver show -subtype sync-source
```

```

node_A_1::> vserver show -subtype sync-source
 Admin Operational Root
Vserver Type Subtype State State Volume
Aggregate

...
vs1a data sync-source
 running running vs1a_vol aggr_b2

```

4. Confirm that the switchback operations succeeded by using the `metrocluster operation show` command.

| If the command output shows...                                                                    | Then...                                                                                                   |
|---------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| That the switchback operation state is successful.                                                | The switchback process is complete and you can proceed with operation of the system.                      |
| That the switchback operation or switchback-continuation-agent operation is partially successful. | Perform the suggested fix provided in the output of the <code>metrocluster operation show</code> command. |

### After you finish

You must repeat the previous sections to perform the switchback in the opposite direction. If site\_A did a switchover of site\_B, have site\_B do a switchover of site\_A.

### Mirroring the root aggregates of the replacement nodes

If disks were replaced, you must mirror the root aggregates of the new nodes on the disaster site.

## Steps

1. On the disaster site, identify the aggregates which are not mirrored:

```
storage aggregate show
```

```
cluster_A::> storage aggregate show

Aggregate Size Available Used% State #Vols Nodes RAID
Status

node_A_1_aggr0
 1.49TB 74.12GB 95% online 1 node_A_1
raid4,

normal
node_A_2_aggr0
 1.49TB 74.12GB 95% online 1 node_A_2
raid4,

normal
node_A_1_aggr1
 1.49TB 74.12GB 95% online 1 node_A_1 raid
4, normal

mirrored
node_A_2_aggr1
 1.49TB 74.12GB 95% online 1 node_A_2 raid
4, normal

mirrored
4 entries were displayed.

cluster_A::>
```

2. Mirror one of the root aggregates:

```
storage aggregate mirror -aggregate root-aggregate
```

The following example shows how the command selects disks and prompts for confirmation when mirroring the aggregate.

```

cluster_A::> storage aggregate mirror -aggregate node_A_2_aggr0

Info: Disks would be added to aggregate "node_A_2_aggr0" on node
"node_A_2" in
 the following manner:

 Second Plex

 RAID Group rg0, 3 disks (block checksum, raid4)
 Position Disk Type
Size

- parity 2.10.0 SSD
894.0GB data 1.11.19 SSD
894.0GB data 2.10.2 SSD

 Aggregate capacity available for volume use would be 1.49TB.

Do you want to continue? {y|n}: y

cluster_A::>

```

3. Verify that mirroring of the root aggregate is complete:

```
storage aggregate show
```

The following example shows that the root aggregates are mirrored.

```

cluster_A::> storage aggregate show

Aggregate Size Available Used% State #Vols Nodes RAID
Status

node_A_1_aggr0
 1.49TB 74.12GB 95% online 1 node_A_1 raid4,
mirrored,
normal

node_A_2_aggr0
 2.24TB 838.5GB 63% online 1 node_A_2 raid4,
mirrored,
normal

node_A_1_aggr1
 1.49TB 74.12GB 95% online 1 node_A_1 raid4,
mirrored,
normal

node_A_2_aggr1
 1.49TB 74.12GB 95% online 1 node_A_2 raid4
mirrored,
normal

4 entries were displayed.

cluster_A::>

```

4. Repeat these steps for the other root aggregates.

Any root aggregate that does not have a status of mirrored must be mirrored.

### Reconfiguring ONTAP Mediator (MetroCluster IP configurations)

If you have a MetroCluster IP configuration that was configured with ONTAP Mediator, you must remove and reconfigure the association with ONTAP Mediator.

#### Before you begin

- You must have the IP address and username and password for ONTAP Mediator.
- ONTAP Mediator must be configured and operating on the Linux host.

#### Steps

1. Remove the existing ONTAP Mediator configuration:

```
metrocluster configuration-settings mediator remove
```

2. Reconfigure the ONTAP Mediator configuration:

```
metrocluster configuration-settings mediator add -mediator-address mediator-
```

IP-address

## Verifying the health of the MetroCluster configuration

You should check the health of the MetroCluster configuration to verify proper operation.

### Steps

1. Check that the MetroCluster is configured and in normal mode on each cluster:

```
metrocluster show
```

```
cluster_A::> metrocluster show
Cluster Entry Name State

Local: cluster_A Configuration state configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-disaster
Remote: cluster_B Configuration state configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-disaster
```

2. Check that mirroring is enabled on each node:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
DR
Group Cluster Node Configuration State DR
Mirroring Mode

1 cluster_A
 node_A_1 configured enabled normal
 cluster_B
 node_B_1 configured enabled normal
2 entries were displayed.
```

3. Check that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A::> metrocluster check run
```

```
Last Checked On: 10/1/2014 16:03:37
```

| Component          | Result |
|--------------------|--------|
| nodes              | ok     |
| lifs               | ok     |
| config-replication | ok     |
| aggregates         | ok     |

4 entries were displayed.

Command completed. Use the `metrocluster check show -instance` command or sub-commands in `metrocluster check` directory for detailed results. To check if the nodes are ready to do a switchover or switchback operation, run `metrocluster switchover -simulate` or `metrocluster switchback -simulate`, respectively.

4. Check that there are no health alerts:

```
system health alert show
```

5. Simulate a switchover operation:

- a. From any node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).

- b. Perform the switchover operation with the `-simulate` parameter:

```
metrocluster switchover -simulate
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

6. For MetroCluster IP configurations using ONTAP Mediator, confirm that ONTAP Mediator is up and operating.

- a. Check that the Mediator disks are visible to the system:

```
storage failover mailbox-disk show
```

The following example shows that the mailbox disks have been recognized.

```

node_A_1::*> storage failover mailbox-disk show
 Mailbox
Node Owner Disk Name Disk UUID

still13-vsim-ucs626g
.
.
 local 0m.i2.3L26
7BBA77C9:AD702D14:831B3E7E:0B0730EE:00000000:00000000:00000000:000000
00:00000000:00000000
 local 0m.i2.3L27
928F79AE:631EA9F9:4DCB5DE6:3402AC48:00000000:00000000:00000000:000000
00:00000000:00000000
 local 0m.i1.0L60
B7BCDB3C:297A4459:318C2748:181565A3:00000000:00000000:00000000:000000
00:00000000:00000000
.
.
.
 partner 0m.i1.0L14
EA71F260:D4DD5F22:E3422387:61D475B2:00000000:00000000:00000000:000000
00:00000000:00000000
 partner 0m.i2.3L64
4460F436:AAE5AB9E:D1ED414E:ABF811F7:00000000:00000000:00000000:000000
00:00000000:00000000
28 entries were displayed.

```

b. Change to the advanced privilege level:

```
set -privilege advanced
```

c. Check that the mailbox LUNs are visible to the system:

```
storage iscsi-initiator show
```

The output will show the presence of the mailbox LUNs:

```

Node Type Label Target Portal Target Name
Admin/Op

.
.
.
.node_A_1
 mailbox
 mediator 172.16.254.1 iqn.2012-
05.local:mailbox.target.db5f02d6-e3d3 up/up
.
.
.
17 entries were displayed.

```

d. Return to the administrative privilege level:

```
set -privilege admin
```

## Recovering from a non-controller failure

After the equipment at the disaster site has undergone any required maintenance or replacement, but no controllers were replaced, you can begin the process of returning the MetroCluster configuration to a fully redundant state. This includes healing the configuration (first the data aggregates and then the root aggregates) and performing the switchback operation.

### Before you begin

- All MetroCluster hardware in the disaster cluster must be functional.
- The overall MetroCluster configuration must be in switchover.
- In a fabric-attached MetroCluster configuration, the ISL must be up and operating between the MetroCluster sites.

### Enable console logging

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

- Leave AutoSupport enabled during maintenance.
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Healing the configuration in a MetroCluster configuration

In MetroCluster FC configurations, you perform the healing operations in a specific order to restore MetroCluster functionality following a switchover.

In MetroCluster IP configurations, healing operations should start automatically following a switchover. If they do not, you can perform the healing operations manually.

### Before you begin

- Switchover must have been performed and the surviving site must be serving data.
- Nodes on the disaster site must be halted or remain powered off.

They must not be fully booted during the healing process.

- Storage at the disaster site must be accessible (shelves are powered up, functional, and accessible).
- In fabric-attached MetroCluster configurations, inter-switch links (ISLs) must be up and operating.
- In four-node MetroCluster configurations, nodes in the surviving site must not be in HA failover state (all nodes must be up and running for each HA pair).

### About this task

The healing operation must first be performed on the data aggregates, and then on the root aggregates.

### Healing the data aggregates

You must heal the data aggregates after repairing and replacing any hardware on the disaster site. This process resynchronizes the data aggregates and prepares the (now repaired) disaster site for normal operation. You must heal the data aggregates prior to healing the root aggregates.

### About this task

The following example shows a forced switchover, where you bring the switched-over aggregate online. All configuration updates in the remote cluster successfully replicate to the local cluster. You power up the storage on the disaster site as part of this procedure, but you do not and must not power up the controller modules on the disaster site.

### Steps

1. Verify that switchover was completed:

```
metrocluster operation show
```

```
controller_A_1::> metrocluster operation show
 Operation: switchover
 State: successful
 Start Time: 7/25/2014 20:01:48
 End Time: 7/25/2014 20:02:14
 Errors: -
```

2. Resynchronize the data aggregates by running the following command from the surviving cluster:

```
metrocluster heal -phase aggregates
```

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `--override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

3. Verify that the operation has been completed:

```
metrocluster operation show
```

```
controller_A_1::> metrocluster operation show
 Operation: heal-aggregates
 State: successful
Start Time: 7/25/2014 18:45:55
 End Time: 7/25/2014 18:45:56
 Errors: -
```

4. Check the state of the aggregates:

```
storage aggregate show command.
```

```
controller_A_1::> storage aggregate show
Aggregate Size Available Used% State #Vols Nodes RAID
Status

...
aggr_b2 227.1GB 227.1GB 0% online 0 mcc1-a2 raid_dp,
mirrored, normal...
```

5. If storage has been replaced at the disaster site, you might need to remirror the aggregates.

## Healing the root aggregates after a disaster

After the data aggregates have been healed, you must heal the root aggregates in preparation for the switchback operation.

### Before you begin

The data aggregates phase of the MetroCluster healing process must have been completed successfully.

### Steps

## 1. Switch back the mirrored aggregates:

```
metrocluster heal -phase root-aggregates
```

```
mccl1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `--override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

## 2. Ensure that the heal operation is complete by running the following command on the destination cluster:

```
metrocluster operation show
```

```
mccl1A::> metrocluster operation show
Operation: heal-root-aggregates
State: successful
Start Time: 7/29/2014 20:54:41
End Time: 7/29/2014 20:54:42
Errors: -
```

## Verifying that your system is ready for a switchback

If your system is already in the switchover state, you can use the `-simulate` option to preview the results of a switchback operation.

### Steps

1. Power up each controller module on the disaster site.

**If the nodes are powered off:**

Power on the nodes.

**If the nodes are at the LOADER prompt:**

Run the command: `boot_ontap`

2. After the node boot completes, verify that the root aggregates are mirrored.

**If a plex fails:**

- a. Destroy the failed plex:

```
storage aggregate plex delete -aggregate <aggregate_name> -plex
<plex_name>
```

- b. Reestablish the mirror relationship by recreating the mirror:

```
storage aggregate mirror -aggregate <aggregate-name>
```

**If a plex is offline:**

Online the plex:

```
storage aggregate plex online -aggregate <aggregate_name> -plex <plex_name>
```

**If both plexes are present:**

Resynchronization starts automatically.

3. Simulate the switchback operation:

- a. From either surviving node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).

- b. Perform the switchback operation with the `-simulate` parameter:

```
metrocluster switchback -simulate
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

4. Review the output that is returned.

The output shows whether the switchback operation would run into errors.

**Example of verification results**

The following example shows the successful verification of a switchback operation:

```

cluster4::*> metrocluster switchback -simulate
(metrocluster switchback)
[Job 130] Setting up the nodes and cluster components for the switchback
operation...DBG:backup_api.c:327:backup_nso_sb_vetocheck : MetroCluster
Switch Back
[Job 130] Job succeeded: Switchback simulation is successful.

cluster4::*> metrocluster op show
(metrocluster operation show)
Operation: switchback-simulate
State: successful
Start Time: 5/15/2014 16:14:34
End Time: 5/15/2014 16:15:04
Errors: -

cluster4::*> job show -name Me*

```

| Job ID | Name                    | Owning Vserver | Node        | State   |
|--------|-------------------------|----------------|-------------|---------|
| 130    | MetroCluster Switchback | cluster4       | cluster4-01 | Success |

```

Description: MetroCluster Switchback Job - Simulation

```

## Performing a switchback

After you heal the MetroCluster configuration, you can perform the MetroCluster switchback operation. The MetroCluster switchback operation returns the configuration to its normal operating state, with the sync-source storage virtual machines (SVMs) on the disaster site active and serving data from the local disk pools.

### Before you begin

- The disaster cluster must have successfully switched over to the surviving cluster.
- Healing must have been performed on the data and root aggregates.
- The surviving cluster nodes must not be in the HA failover state (all nodes must be up and running for each HA pair).
- The disaster site controller modules must be completely booted and not in the HA takeover mode.
- The root aggregate must be mirrored.
- The Inter-Switch Links (ISLs) must be online.
- Any required licenses must be installed on the system.

### Steps

1. Confirm that all nodes are in the enabled state:

```
metrocluster node show
```

The following example displays the nodes that are in the "enabled" state:

```
cluster_B::> metrocluster node show

DR
Group Cluster Node Configuration DR

1 cluster_A
 node_A_1 configured enabled heal roots completed
 node_A_2 configured enabled heal roots completed
 cluster_B
 node_B_1 configured enabled waiting for
switchback recovery
 node_B_2 configured enabled waiting for
switchback recovery
4 entries were displayed.
```

2. Confirm that resynchronization is complete on all SVMs:

```
metrocluster vserver show
```

3. Verify that any automatic LIF migrations being performed by the healing operations have been successfully completed:

```
metrocluster check lif show
```

4. Perform the switchback by running the following command from any node in the surviving cluster.

```
metrocluster switchback
```

5. Check the progress of the switchback operation:

```
metrocluster show
```

The switchback operation is still in progress when the output displays "waiting-for-switchback":

```
cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
Mode switchover
AUSO Failure Domain -
Remote: cluster_A Configuration state configured
Mode waiting-for-switchback
AUSO Failure Domain -
```

The switchback operation is complete when the output displays "normal":

```

cluster_B::> metrocluster show
Cluster Entry Name State

Local: cluster_B Configuration state configured
 Mode normal
 AUSO Failure Domain -
Remote: cluster_A Configuration state configured
 Mode normal
 AUSO Failure Domain -

```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the following command at the advanced privilege level.

```
metrocluster config-replication resync-status show
```

#### 6. Reestablish any SnapMirror or SnapVault configurations.

In ONTAP 8.3, you need to manually reestablish a lost SnapMirror configuration after a MetroCluster switchback operation. In ONTAP 9.0 and later, the relationship is reestablished automatically.

## Verifying a successful switchback

After performing the switchback, you want to confirm that all aggregates and storage virtual machines (SVMs) are switched back and online.

### Steps

1. Verify that the switched-over data aggregates are switched back:

```
storage aggregate show
```

In the following example, aggr\_b2 on node B2 has switched back:

```

node_B_1::> storage aggregate show
Aggregate Size Available Used% State #Vols Nodes RAID
Status

...
aggr_b2 227.1GB 227.1GB 0% online 0 node_B_2 raid_dp,
mirrored,
normal

node_A_1::> aggr show
Aggregate Size Available Used% State #Vols Nodes RAID
Status

...
aggr_b2 - - - unknown - node_A_1

```

If the disaster site included unmirrored aggregates and the unmirrored aggregates are no longer present, the aggregate might show up with a state of "unknown" in the output of the `storage aggregate show` command. Contact technical support to remove the out-of-date entries for the unmirrored aggregates and reference the Knowledge Base article [How to remove stale unmirrored aggregate entries in a MetroCluster following disaster where storage was lost](#).

2. Verify that all sync-destination SVMs on the surviving cluster are dormant (showing an operational state of "stopped"):

```
vserver show -subtype sync-destination
```

```

node_B_1::> vserver show -subtype sync-destination
Vserver Type Subtype Admin Operational Root
Aggregate State State Volume

...
cluster_A-vs1a-mc data sync-destination
 running stopped vs1a_vol aggr_b2

```

Sync-destination aggregates in the MetroCluster configuration have the suffix "-mc" automatically appended to their name to help identify them.

3. Verify the sync-source SVMs on the disaster cluster are up and running:

```
vserver show -subtype sync-source
```

```
node_A_1::> vserver show -subtype sync-source
Vserver Type Subtype Admin Operational Root
Aggregate

...
vs1a data sync-source running running vs1a_vol aggr_b2
```

4. Confirm that the switchback operations succeeded:

```
metrocluster operation show
```

| If the command output shows...                                                                                 | Then...                                                                                                   |
|----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| That the switchback operation state is successful.                                                             | The switchback process is complete and you can proceed with operation of the system.                      |
| That the switchback operation or <code>switchback-continuation-agent</code> operation is partially successful. | Perform the suggested fix provided in the output of the <code>metrocluster operation show</code> command. |

### After you finish

You must repeat the previous sections to perform the switchback in the opposite direction. If site\_A did a switchover of site\_B, have site\_B do a switchover of site\_A.

## Deleting stale aggregate listings after switchback

In some circumstances after switchback, you might notice the presence of *stale* aggregates. Stale aggregates are aggregates that have been removed from ONTAP, but whose information remains recorded on disk. Stale aggregates are displayed with the `nodeshell aggr status -r` command but not with the `storage aggregate show` command. You can delete these records so that they no longer appear.

### About this task

Stale aggregates can occur if you relocated aggregates while the MetroCluster configuration was in switchover. For example:

1. Site A switches over to Site B.
2. You delete the mirroring for an aggregate and relocate the aggregate from node\_B\_1 to node\_B\_2 for load balancing.
3. You perform aggregate healing.

At this point a stale aggregate appears on node\_B\_1, even though the actual aggregate has been deleted from that node. This aggregate appears in the output from the `nodeshell aggr status -r` command. It does not appear in the output of the `storage aggregate show` command.

1. Compare the output of the following commands:

```
storage aggregate show
```

```
run local aggr status -r
```

Stale aggregates appear in the `run local aggr status -r` output but not in the `storage aggregate show` output. For example, the following aggregate might appear in the `run local aggr status -r` output:

```
Aggregate aggr05 (failed, raid_dp, partial) (block checksums)
Plex /aggr05/plex0 (offline, failed, inactive)
 RAID group /myaggr/plex0/rg0 (partial, block checksums)

 RAID Disk Device HA SHELF BAY CHAN Pool Type RPM Used (MB/blks)
 Phys (MB/blks)

 dparity FAILED N/A 82/ -
 parity 0b.5 0b - - SA:A 0 VMDISK N/A 82/169472
88/182040
 data FAILED N/A 82/ -
 data FAILED N/A 82/ -
Raid group is missing 7 disks.
```

2. Remove the stale aggregate:

- a. From either node's prompt, change to the advanced privilege level:

```
set -privilege advanced
```

You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).

- b. Remove the stale aggregate:

```
aggregate remove-stale-record -aggregate aggregate_name
```

- c. Return to the admin privilege level:

```
set -privilege admin
```

3. Confirm that the stale aggregate record was removed:

```
run local aggr status -r
```

# Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

## Copyright

<https://www.netapp.com/company/legal/copyright/>

## Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

## Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

## Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

## Safety information and regulatory notices

[https://library.netapp.com/ecm/ecm\\_download\\_file/ECMP12475945](https://library.netapp.com/ecm/ecm_download_file/ECMP12475945)

## Copyright information

Copyright © 2026 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.