



# **Upgrade, refresh, or expand the MetroCluster configuration**

## **ONTAP MetroCluster**

NetApp  
February 13, 2026

This PDF was generated from [https://docs.netapp.com/us-en/ontap-metrocluster/upgrade/concept\\_choosing\\_an\\_upgrade\\_method\\_mcc.html](https://docs.netapp.com/us-en/ontap-metrocluster/upgrade/concept_choosing_an_upgrade_method_mcc.html) on February 13, 2026. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Upgrade, refresh, or expand the MetroCluster configuration .....	1
Start here - Choose your procedure .....	1
Start here: Choose between controller upgrade, system refresh, or expansion .....	1
Choose a controller upgrade procedure .....	2
Choosing a system refresh method .....	8
Choose an expansion procedure .....	12
Upgrade controllers in four-node MetroCluster IP using switchover and switchback with "system controller replace" commands (ONTAP 9.13.1 and later) .....	13
Workflow for upgrading MetroCluster IP controllers using "system controller replace" commands (ONTAP 9.13.1 or later) .....	13
Prepare to upgrade .....	14
Upgrade your controllers .....	26
Complete the MetroCluster IP controller upgrade .....	51
Upgrade controllers in MetroCluster IP using switchover and switchback (ONTAP 9.8 and later) .....	52
Workflow for MetroCluster IP controller upgrades using switchover and switchback (ONTAP 9.8 and later) .....	52
Prepare to upgrade .....	53
Upgrade your controllers .....	67
Complete the MetroCluster IP controller upgrade .....	95
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) .....	97
Enable console logging .....	97
Workflow for upgrading controllers in a MetroCluster IP configuration .....	98
Prepare for the upgrade .....	98
Switch over the MetroCluster configuration .....	104
Remove AFF A700 or FAS9000 platform controller module and NVS .....	105
Install the AFF A900 or FAS9500 NVS and controller modules .....	107
Update the switch RCF files to accommodate the new platforms .....	113
Configure the new controllers .....	114
Switch back the MetroCluster configuration .....	122
Check the health of the MetroCluster configuration .....	124
Upgrade the nodes on site_A .....	125
Restore Tiebreaker or Mediator monitoring .....	125
Send a custom AutoSupport message after maintenance .....	125
Upgrade controllers in a MetroCluster FC configuration using switchover and switchback .....	126
Supported platform combinations .....	126
About this task .....	126
Enable console logging .....	127
Prepare for the upgrade .....	128
Switch over the MetroCluster configuration .....	134
Prepare the network configuration of the old controllers .....	136
Remove the old platforms .....	138
Configure the new controllers .....	138

Switch back the MetroCluster configuration . . . . .	149
Check the health of the MetroCluster configuration . . . . .	151
Upgrade the nodes on cluster_A . . . . .	152
Send a custom AutoSupport message after maintenance . . . . .	152
Restore Tiebreaker monitoring . . . . .	152
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) . . . . .	152
Enable console logging . . . . .	153
Prepare for the upgrade . . . . .	154
Clear slot 7 on the AFF A700 controller . . . . .	154
Switch over the MetroCluster configuration . . . . .	160
Remove the AFF A700 or FAS9000 controller module and NVS at site_B . . . . .	162
Remove the controller module and NVS from both nodes at site_B . . . . .	162
Install the AFF A900 or FAS9500 NVS and controller module . . . . .	164
Switch back the MetroCluster configuration . . . . .	177
Check the health of the MetroCluster configuration . . . . .	178
Upgrade the nodes on site_A . . . . .	179
Send a custom AutoSupport message after maintenance . . . . .	179
Restore Tiebreaker monitoring . . . . .	179
Upgrade controllers in a four-node MetroCluster FC configuration using switchover and switchback with "system controller replace" commands (ONTAP 9.10.1 and later) . . . . .	180
Supported platform combinations . . . . .	180
About this task . . . . .	180
Enable console logging . . . . .	181
Prepare for the upgrade . . . . .	181
Replace the old controllers and boot up the new controllers . . . . .	185
Complete the upgrade . . . . .	196
Upgrade the nodes on cluster_A . . . . .	197
Refreshing a four-node MetroCluster FC configuration . . . . .	197
Enable console logging . . . . .	197
Perform the refresh procedure . . . . .	197
Refresh a four-node or an eight-node MetroCluster IP configuration (ONTAP 9.8 and later) . . . . .	199
Enable console logging . . . . .	201
Perform the refresh procedure . . . . .	201
Expand a two-node MetroCluster FC configuration to a four-node configuration . . . . .	210
Expanding a two-node MetroCluster FC configuration to a four-node configuration . . . . .	210
Enable console logging . . . . .	212
Verifying the state of the MetroCluster configuration . . . . .	212
Sending a custom AutoSupport message before adding nodes to the MetroCluster configuration . . . . .	214
Zoning for the new controller ports when adding a controller module in a fabric-attached MetroCluster configuration . . . . .	214
Add a new controller module to each cluster . . . . .	215
Refreshing the MetroCluster configuration with new controllers . . . . .	246
Enabling storage failover on both controller modules and enabling cluster HA . . . . .	249
Restarting the SVMs . . . . .	250

Expand a four-node MetroCluster FC configuration to an eight-node configuration . . . . .	250
Expanding a four-node MetroCluster FC configuration to an eight-node configuration . . . . .	250
Enable console logging . . . . .	253
Determining the new cabling layout . . . . .	253
Racking the new equipment . . . . .	254
Verifying the health of the MetroCluster configuration . . . . .	254
Checking for MetroCluster configuration errors with Config Advisor . . . . .	256
Sending a custom AutoSupport message prior to adding nodes to the MetroCluster configuration . . . . .	256
Recable and zone a switch fabric for the new nodes . . . . .	256
Configure ONTAP on the new controllers . . . . .	258
Checking for MetroCluster configuration errors with Config Advisor . . . . .	288
Sending a custom AutoSupport message after to adding nodes to the MetroCluster configuration . . . . .	288
Verifying switchover, healing, and switchback . . . . .	289
Expand a MetroCluster IP configuration . . . . .	289
Enable console logging . . . . .	291
Example naming in this procedure . . . . .	291
Supported platform combinations when adding a second DR group . . . . .	291
Sending a custom AutoSupport message prior to maintenance . . . . .	294
Considerations for VLANs when adding a new DR group . . . . .	295
Verifying the health of the MetroCluster configuration . . . . .	296
Removing the configuration from monitoring applications . . . . .	300
Preparing the new controller modules . . . . .	301
Upgrade RCF files . . . . .	301
Join the new nodes to the clusters . . . . .	303
Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates . . . . .	304
Finalizing the addition of the new nodes . . . . .	316
Remove a DR group from a MetroCluster configuration . . . . .	322
Enable console logging . . . . .	322
Remove the DR group nodes from each cluster . . . . .	323
Where to find additional information . . . . .	329
MetroCluster and miscellaneous information . . . . .	329

# 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	<a href="#">Choose a controller upgrade procedure</a>
System refresh	<a href="#">Choose a system refresh procedure</a>
Expansion	<ul style="list-style-type: none"><li>• Two-node MetroCluster to four</li><li>• Four-node MetroCluster FC to eight</li><li>• Four-node MetroCluster IP to eight</li></ul>

## 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
Source MetroCluster IP platform	AFF A150		Note 5							
	AFF A20									
	FAS2750 AFF A220		Note 5							
	FAS500f AFF C250 AFF A250						Note 4			
	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 AFF A300					Note 3		Note 2			Note 3	
	AFF A320											
	FAS8300 AFF C400 AFF A400					Note 3		Note 2	Note 3		Note 3	
	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
	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
	FAS8300									Note 4	
	AFF A400									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?
ASA A400	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	<a href="#">Link to procedure</a>
FC	Upgrade with 'system controller replace' commands	9.10.1 and later	<a href="#">Link to procedure</a>
FC	Manual upgrade with CLI commands (AFF A700/FAS9000 to AFF A900/FAS9500 only)	9.10.1 and later	<a href="#">Link to procedure</a>
IP	Manual upgrade with CLI commands (AFF A700/FAS9000 to AFF A900/FAS9500 only)	9.10.1 and later	<a href="#">Link to procedure</a>
FC	Manual upgrade with CLI commands	9.8 and later	<a href="#">Link to procedure</a>
IP	Manual upgrade with CLI commands	9.8 and later	<a href="#">Link to procedure</a>

## 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	<a href="#">Link to procedure</a>
FC	Using system controller replace commands	9.8 and later	<a href="#">Link to procedure</a>
FC	Using system controller replace commands	9.5 through 9.7	<a href="#">Link to procedure</a>
FC	Using manual ARL commands	9.8	<a href="#">Link to procedure</a>
FC	Using manual ARL commands	9.7 and earlier	<a href="#">Link to procedure</a>

## 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				Note 1		Note 1	Note 1			
	AFF A220										
	FAS500f				Note 1		Note 1	Note 1			
	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										

### 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 then 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	<a href="#">Link to procedure</a>
• Method: Expand the MetroCluster configuration and then remove the old nodes	Four-node IP	9.8 and later	<a href="#">Link to procedure</a>

### 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)	<a href="#">Link to procedure</a>
Method: Expand a four-node MetroCluster FC to eight	Four-node FC	ONTAP 9 or later	<a href="#">Link to procedure</a>
Method: Expand a four-node MetroCluster IP to eight	Four-node IP	ONTAP 9.9.1 and later	<a href="#">Link to procedure</a>

## 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.

### 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
Source MetroCluster IP platform	AFF A150										
	AFF A20										
	FAS2750										
	AFF A220										
	FAS500f										
	AFF C250										
	AFF A250										
	FAS50										
	AFF C30										
	AFF A20										
	FAS8200					Note 2					
	AFF A300										
	AFF A320										
	FAS8300					Note 2					
	AFF C400										
	AFF A400										
	FAS8700										
	AFF C60										
	AFF A50										
	FAS70										
	FAS9000					Note 2					
	AFF A700										
	AFF A70										
	AFF C800										
	AFF A800										
	FAS9500										
	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>
<b>Note:</b> 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**

- 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**

- At the LOADER prompt of the halted node, enter the following:

```
setenv bootarg.siw.interop_enabled 1  
saveenv  
printenv bootarg.siw.interop_enabled
```

- Boot the node:

```
boot_ontap
```

- When the node boots, perform a giveback for the node at the prompt:

```
storage failover giveback -ofnode <node_name>
```

- 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          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 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
- vserver 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
- 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 [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"><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>

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	<a href="#">Substep b</a>
Use a different home port	<a href="#">Step 4</a>

- 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 data1if1 -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 -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>
```

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

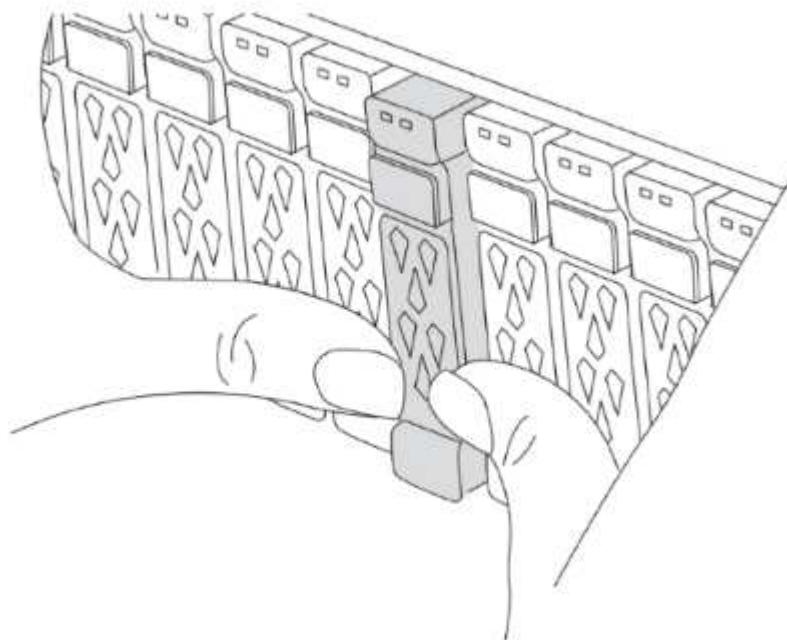
1. Plan out the positioning of the new controller modules and storage shelves as needed.

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"><li>• An AFF A150 to an AFF A20 system</li><li>• An AFF A220 to an AFF A20 system</li></ul>	Controller to external shelf conversion
<ul style="list-style-type: none"><li>• An AFF A800 to an AFF A90 system</li><li>• An AFF C800 to an AFF C80 system</li><li>• An AFF A250 to an AFF A30 system</li><li>• An AFF C250 to an AFF C30 system</li><li>• An AFF A70 to an AFF A90 system</li></ul>	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.

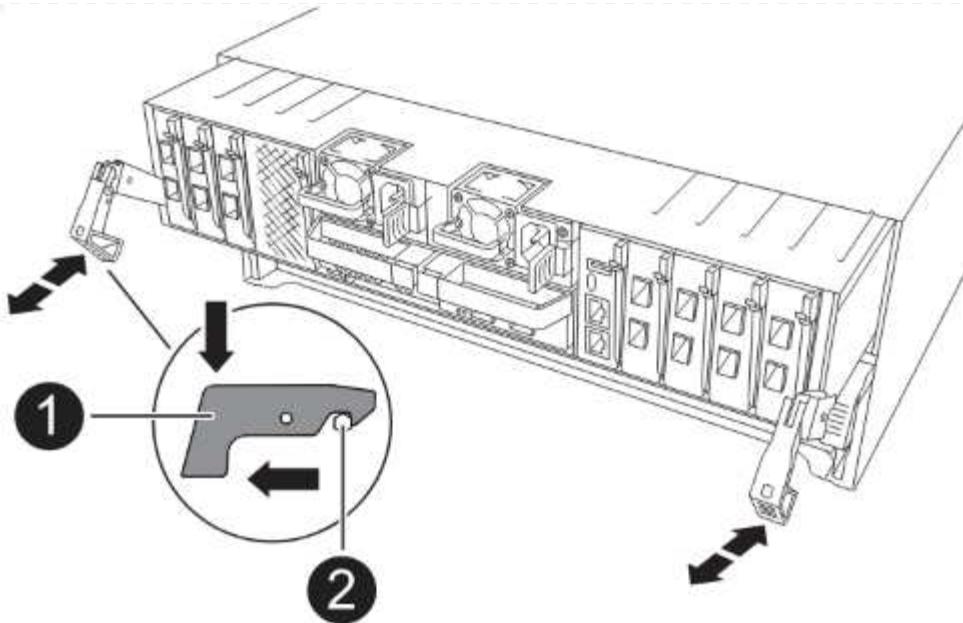
Each drive shelf, including the chassis, requires a unique ID.

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.



#### All other upgrades

Install the controller modules in the rack or cabinet.

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	<code>ucadmin modify -m fc -t initiator &lt;adapter-name&gt;</code>

CNA Ethernet	ucadmin modify -mode cna <adapter-name>
FC target	fcadmin config -t target <adapter-name>
FC initiator	fcadmin config -t initiator <adapter-name>

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	ucadmin show
FC	fcadmin show

**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"> <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. 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
```

3. 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
          Usable
Disk      Size      Disk      Container  Container
Owner
          Shelf Bay Type    Type    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          894.0GB    11     0  SSD      shared    testaggr
node_A_1
1.11.1          894.0GB    11     1  SSD      shared    testaggr
node_A_1
1.11.2          894.0GB    11     2  SSD      shared    testaggr
node_A_1
```

4. 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
```

5. 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`.

a. 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    S0M1J2CF
node_B_1-old(4068741254)  node_B_1-old(4068741254)
rr18:8.126L3  node_B_1-old(4068741254)    Pool0    S0M0CQM5
node_B_1-old(4068741254)  node_B_1-old(4068741254)
rr18:9.126L27 node_B_1-old(4068741254)    Pool0    S0M1PSDW
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  S0M1J2CF
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L29  node_B_1-new(1574774970)  Pool0  S0M0CQM5
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L1   node_B_1-new(1574774970)  Pool0  S0M1PSDW
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:

- View the broadcast domains:

```
network port broadcast-domain show
```

- 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>
```

- Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

- Add the physical port that will host the intercluster LIFs to the corresponding broadcast domain.

- Modify intercluster LIFs to use the new physical port as home port.

- 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](#)

- 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](#)

- 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 <a href="#">Restore onboard key management encryption keys</a>.</p>

<p>External key management</p>	<pre>security key-manager external restore -vserver &lt;svm-name&gt; -node &lt;node-name&gt; -key-server &lt;host_name IP_address:port&gt; -key-id &lt;key_id&gt; -key-tag key_tag &lt;node-name&gt;</pre> <p>For more information, see <a href="#">Restore external key management encryption keys</a>.</p>
--------------------------------	--

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>
<b>Note:</b> 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**

- 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**

- At the LOADER prompt of the halted node, enter the following:

```
setenv bootarg.siw.interop_enabled 1  
saveenv  
printenv bootarg.siw.interop_enabled
```

- Boot the node:

```
boot_ontap
```

- When the node boots, perform a giveback for the node at the prompt:

```
storage failover giveback -ofnode <node_name>
```

- 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.

	node_A_1-old			node_A_1-new		
LIF	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:

- 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 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	<a href="#">Removing MetroCluster Configurations</a>
Mediator	Issue the following command from the ONTAP prompt:  <code>metrocluster configuration-settings mediator remove</code>
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

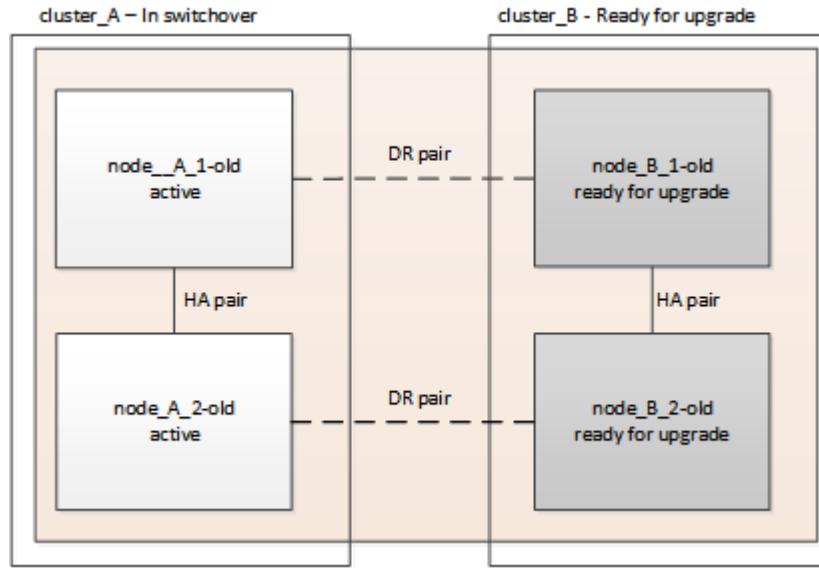
You switch over the configuration to site\_A so that the platforms on site\_B can be upgraded.

#### 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"><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>

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	<a href="#">Substep b</a>
Use a different home port	<a href="#">Step 4</a>

- 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 data1if1 -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 a1a -port e0d
```

c. Remove VLAN and interface group ports from the broadcast domain:

```
network port broadcast-domain remove-ports -ipspace <ipspace> -broadcast
-domain <broadcast-domain-name> -ports
<nodename:portname, nodename:portnamee>,...
```

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

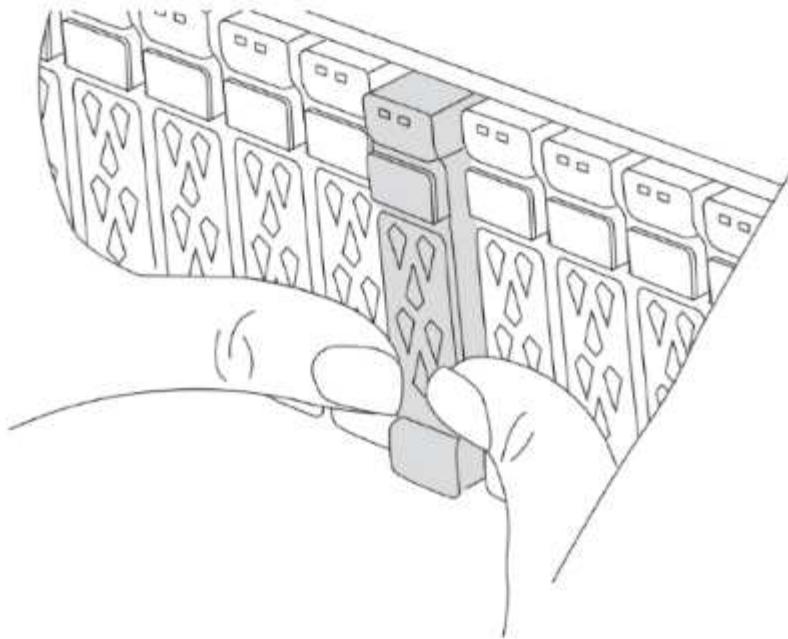
1. Plan out the positioning of the new controller modules and storage shelves as needed.

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"><li>• An AFF A150 to an AFF A20 system</li><li>• An AFF A220 to an AFF A20 system</li></ul>	Controller to external shelf conversion
<ul style="list-style-type: none"><li>• An AFF A800 to an AFF A90 system</li><li>• An AFF C800 to an AFF C80 system</li><li>• An AFF A250 to an AFF A30 system</li><li>• An AFF C250 to an AFF C30 system</li><li>• An AFF A70 to an AFF A90 system</li></ul>	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.

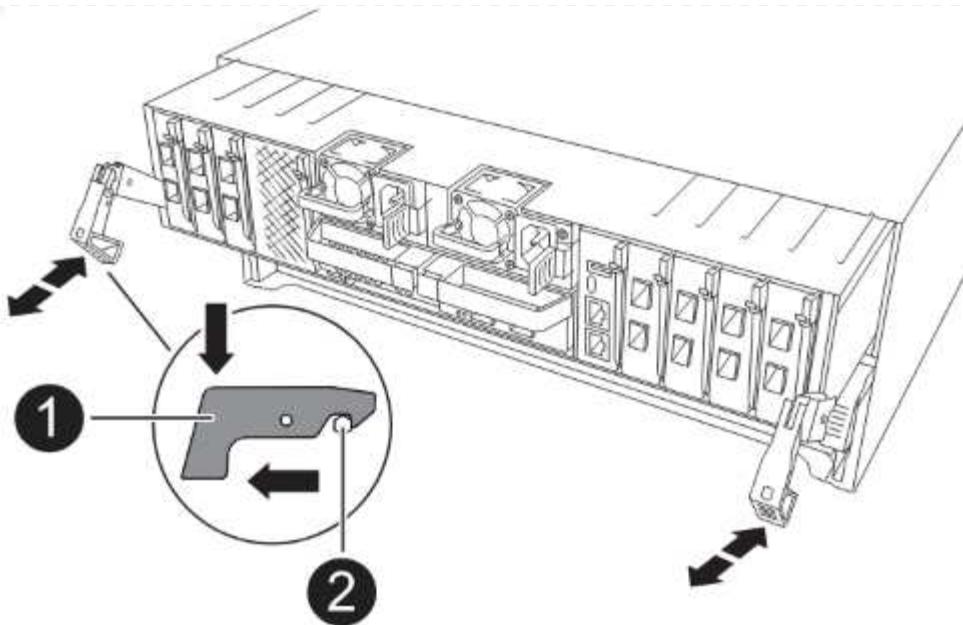
Each drive shelf, including the chassis, requires a unique ID.

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.



#### All other upgrades

Install the controller modules in the rack or cabinet.

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.

<b>If you have this type of HBA and desired mode...</b>	<b>Use this command...</b>
---	----------------------------

CNA FC	ucadmin modify -m fc -t initiator <adapter-name>
CNA Ethernet	ucadmin modify -mode cna <adapter-name>
FC target	fcadmin config -t target <adapter-name>
FC initiator	fcadmin config -t initiator <adapter-name>

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	ucadmin show
FC	fcadmin show

**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"><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. 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
```

3. 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
          Usable
Disk       Size      Disk      Container  Container
Owner
          Shelf Bay Type    Type     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          894.0GB   11     0   SSD    shared    testaggr
node_A_1
1.11.1          894.0GB   11     1   SSD    shared    testaggr
node_A_1
1.11.2          894.0GB   11     2   SSD    shared    testaggr
node_A_1
```

4. 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
```

5. 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`.

a. 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/switchback 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)  Pool0  S0M1J2CF
node_B_1-old(537403322)  node_B_1-old(537403322)
prod2-rk18:8.126L3   node_B_1-old(537403322)  Pool0  S0M0CQM5
node_B_1-old(537403322)  node_B_1-old(537403322)
prod1-rk18:9.126L27  node_B_1-old(537403322)  Pool0  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	disk reassign -s <old-sysid> -d <new-sysid> -r <dr-partner-sysid>
No	disk reassign -s <old-sysid> -d <new-sysid>

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)  Pool0  S0M1J2CF
node_B_1-new(537097247)  node_B_1-new(537097247)
prod02-rk18:9.126L29 node_B_1-new(537097247)  Pool0  S0M0CQM5
node_B_1-new(537097247)  node_B_1-new(537097247)
prod01-rk18:8.126L1  node_B_1-new(537097247)  Pool0  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
                                         fast zeroed
                                         64-bit
                                         mirror_resync_priority=high(fixed)

```

## 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	<code>security key-manager onboard sync</code>  For more information, see <a href="#">Restore onboard key management encryption keys</a> .
External key management	<code>security key-manager external restore -vserver &lt;SVM&gt; -node &lt;node&gt; -key -server &lt;host_name IP_address:port&gt; -key-id key_id -key-tag key_tag &lt;node_name&gt;</code>  For more information, see <a href="#">Restore external key management encryption keys</a> .

12. Verify that all ports are in a broadcast domain:

- 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](#)

[Combining 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>
```

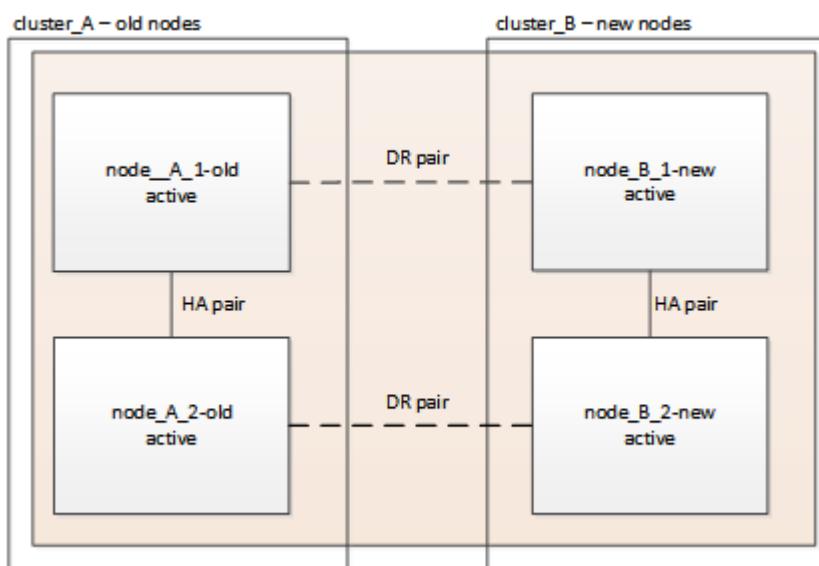
Perform this step on all SVMs as required.

### 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	<a href="#">Add MetroCluster configurations.</a>
Mediator	<a href="#">Configure ONTAP Mediator from a MetroCluster IP configuration.</a>
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.

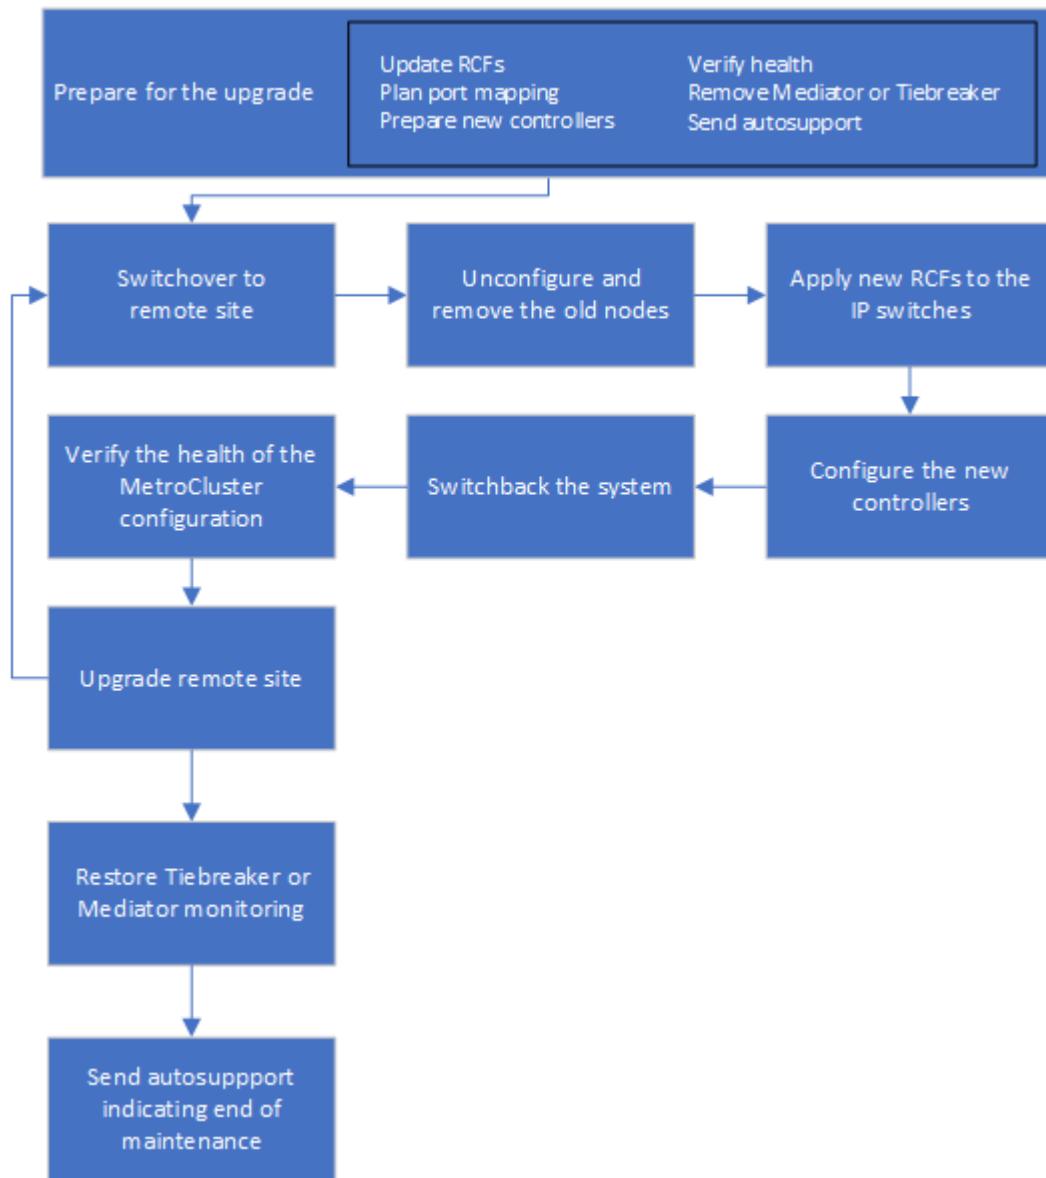
- Trigger a maintenance AutoSupport message before and after maintenance to disable case creation for the duration of the maintenance activity.

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
• AFF A900 or FAS9500	<ul style="list-style-type: none"><li>• 10</li><li>• 20</li><li>• Any value in the range 101 to 4096 inclusive.</li></ul>

- 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

- 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](#).
  - If Onboard Key Manager is configured: `security key-manager onboard show-backup`  
You will need the passphrase later in the upgrade procedure.
  - If enterprise key management (KMIP) is configured, issue the following commands:

```
security key-manager external show -instance
security key-manager key query
```

- 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

- Collect the output for the following command:

```
storage iscsi-initiator show
```

- 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 in the MetroCluster Tiebreaker Installation and Configuration content</a>
Mediator	<p>Issue the following command from the ONTAP prompt:</p> <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.

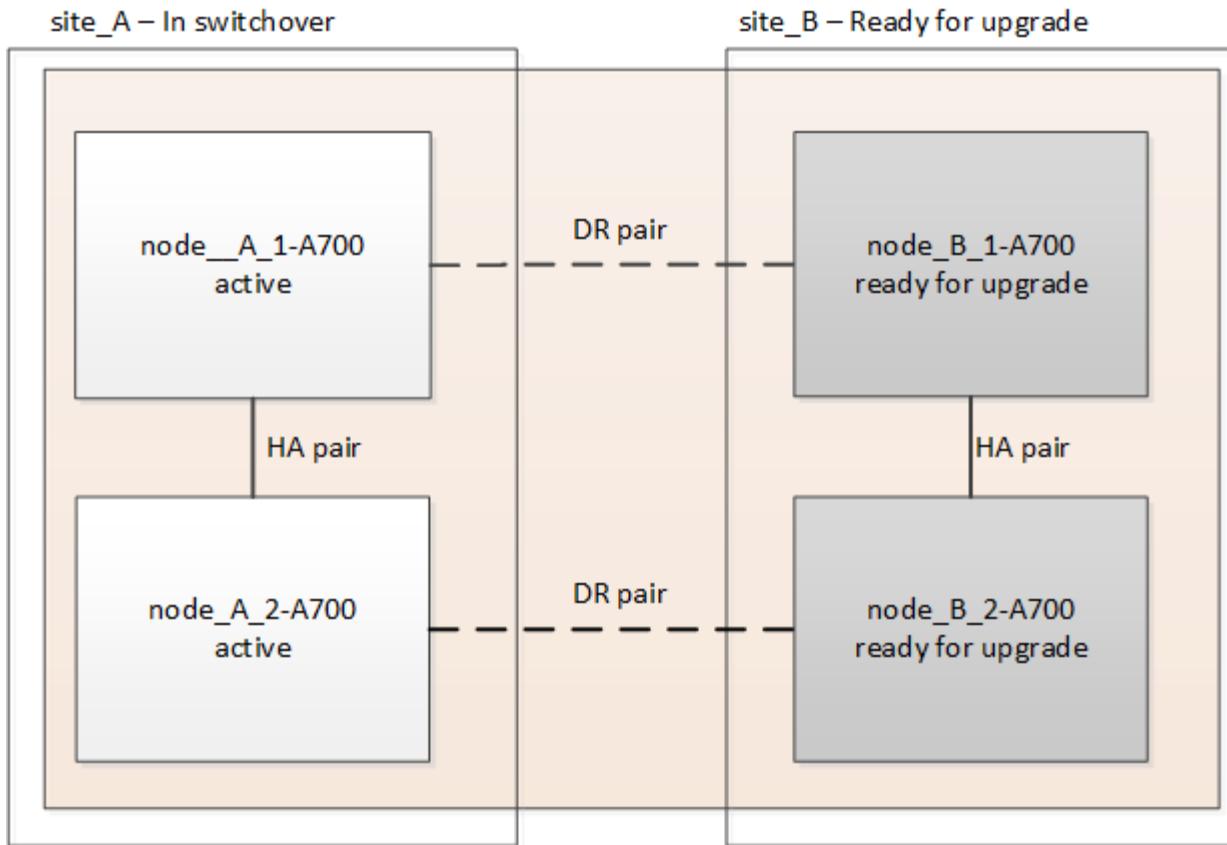
## Switch over the MetroCluster configuration

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

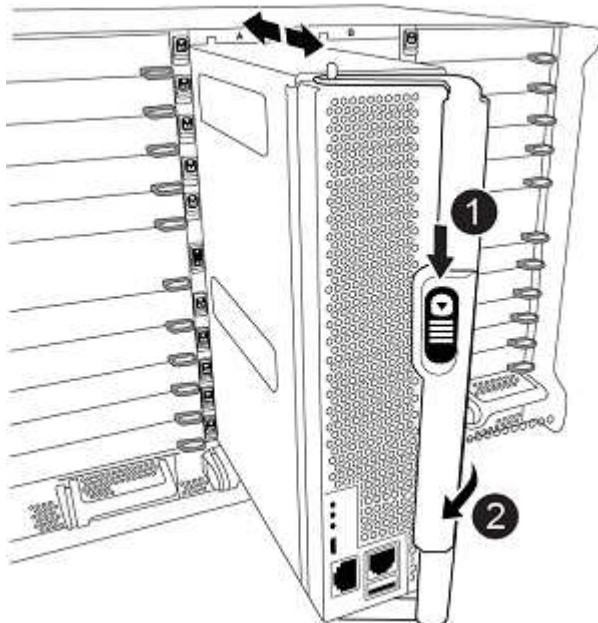
1. Gather the bootarg values from both nodes at site\_B: `printenv`
2. Power off the chassis at site\_B.

## 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

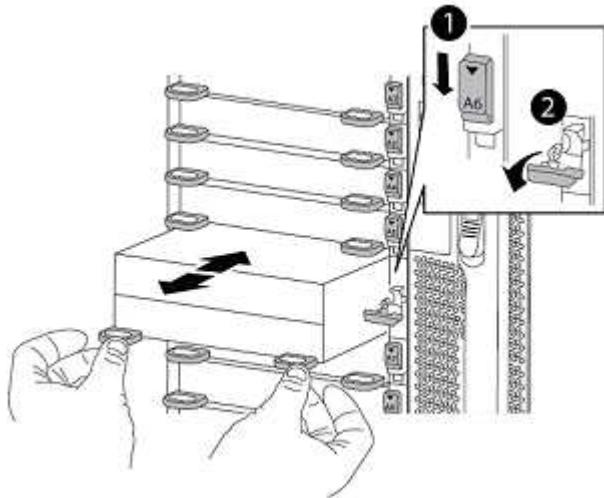
Use the following procedure to remove the AFF A700 or FAS9000 NVS module.

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

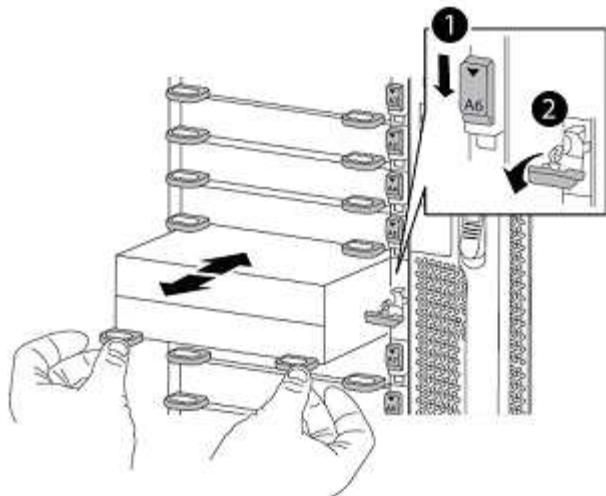
If you are not already grounded, properly ground yourself.

### 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.



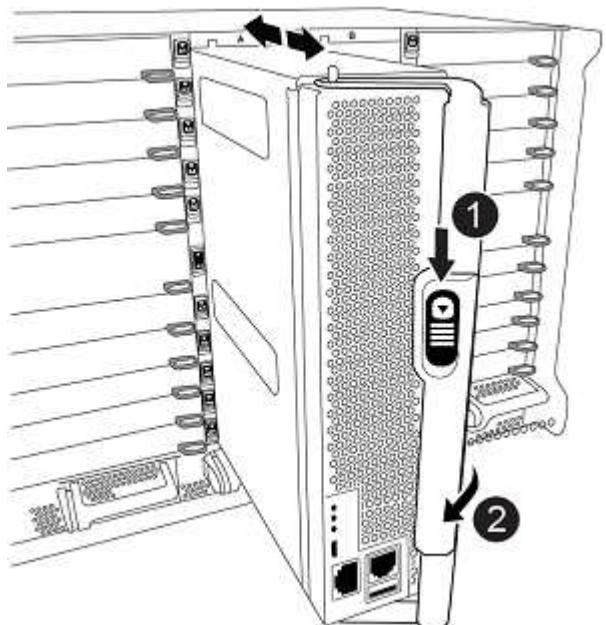
	Lettered and numbered I/O cam latch
	I/O latch completely unlocked

### 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.

- Move the e5b connection to e7b.
- 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

- Access the [NetApp Support Site](#) to download the files used for performing the netboot of the system.
- 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: <code>ifconfig e0M -addr=&lt;filer_addr&gt; -mask=&lt;netmask&gt; -gw=&lt;gateway&gt; -dns=&lt;dns_addr&gt; domain=&lt;dns_domain&gt;</code>  <code>&lt;filer_addr&gt;</code> is the IP address of the storage system. <code>&lt;netmask&gt;</code> is the network mask of the storage system. <code>&lt;gateway&gt;</code> is the gateway for the storage system. <code>&lt;dns_addr&gt;</code> is the IP address of a name server on your network. This parameter is optional. <code>&lt;dns_domain&gt;</code> 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.

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 <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 `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)  Pool0  S0M1J2CF
node_B_1-(537403322)  node_B_1-A700(537403322)
prod2-rk18:8.126L3    node_B_1-A700(537403322)  Pool0  S0M0CQM5
node_B_1-A700(537403322)  node_B_1-A700(537403322)
prod1-rk18:9.126L27   node_B_1-A700(537403322)  Pool0  S0M1PSDW
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	disk reassign -s <i>old-sysid</i> -d <i>new-sysid</i> -r <i>dr-partner-sysid</i>
No	disk reassign -s <i>old-sysid</i> -d <i>new-sysid</i>

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)  Pool0    S0M1J2CF
node_B_1-A900(537097247)  node_B_1-A900(537097247)
prod02-rk18:9.126L29 node_B_1-A900(537097247)  Pool0    S0M0CQM5
node_B_1-A900(537097247)  node_B_1-A900(537097247)
prod01-rk18:8.126L1  node_B_1-A900(537097247)  Pool0    S0M1PSDW
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	<code>security key-manager onboard sync</code>  For more information, see <a href="#">Restoring onboard key management encryption keys</a> .
External key management	<code>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</code>  For more information, see <a href="#">Restoring external key management encryption keys</a> .

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
```

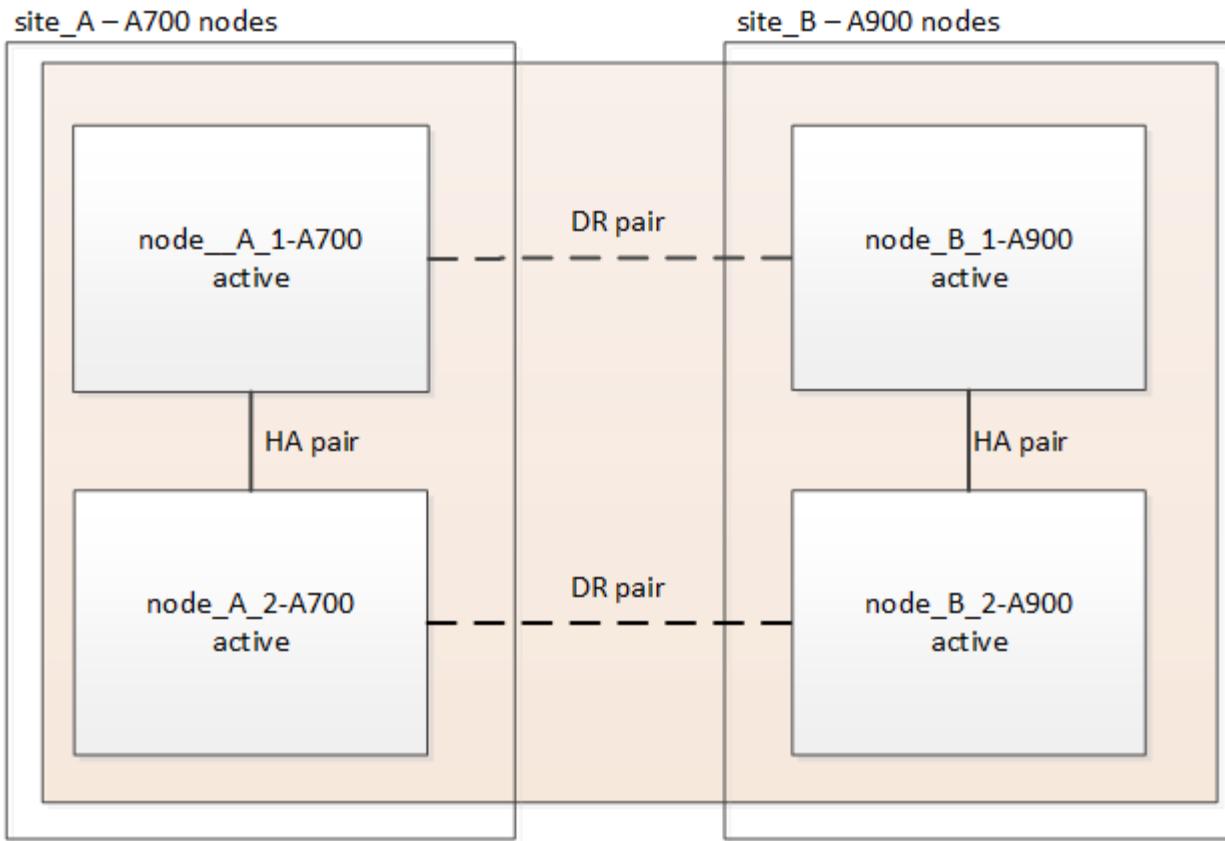
2. Revert the interfaces to their home node:

```
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	<a href="#">Adding MetroCluster configurations</a> in the <i>MetroCluster Tiebreaker Installation and Configuration</i> section.
Mediator	<a href="#">Configure ONTAP Mediator from a MetroCluster IP configuration</a> 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. <a href="#">Prepare cluster ports on an existing controller module</a>

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.

	node_A_1-old			node_A_1-new		
LIF	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:

- 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`

**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.

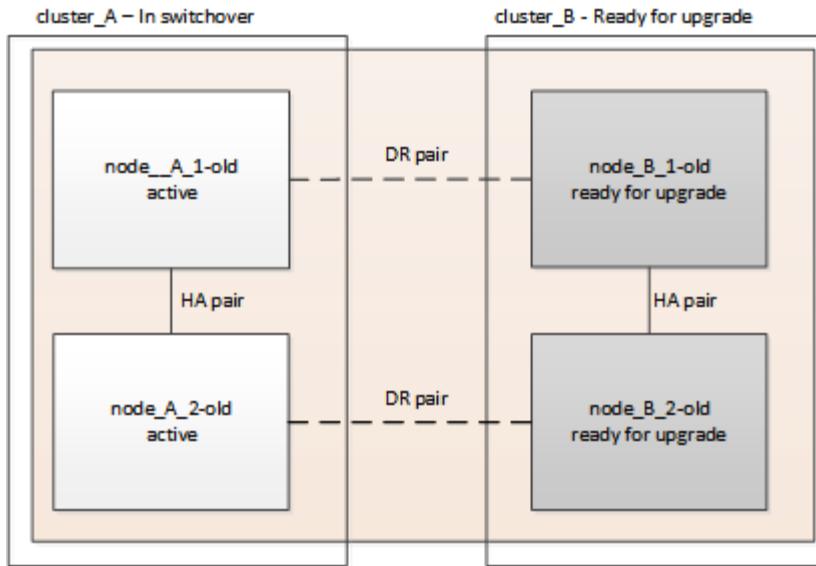
### **Switch over the MetroCluster configuration**

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 data1if1 -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 a1a -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 node-name:portname, node-name: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 ontap-version_image.tgzfile to the target directory: tar -zvxf ontap-version_image.tgz</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	ucadmin show
FC	fcadmin show

### 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</pre> <pre>ha-config modify chassis mcc-2n</pre>
<b>Four or eight nodes</b>	<pre>ha-config modify controller mcc</pre> <pre>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 [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)  Pool0    S0M1J2CF
node_B_1-old(4068741254)  node_B_1-old(4068741254)
rr18:8.126L3  node_B_1-old(4068741254)  Pool0    S0M0CQM5
node_B_1-old(4068741254)  node_B_1-old(4068741254)
rr18:9.126L27 node_B_1-old(4068741254)  Pool0    S0M1PSDW
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  S0M1J2CF
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L29  node_B_1-new(1574774970)  Pool0  S0M0CQM5
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L1   node_B_1-new(1574774970)  Pool0  S0M1PSDW
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	recover_onboard_keymanager
External key management	recover_external_keymanager

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	recover_onboard_keymanager
External key management	recover_external_keymanager

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 <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 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

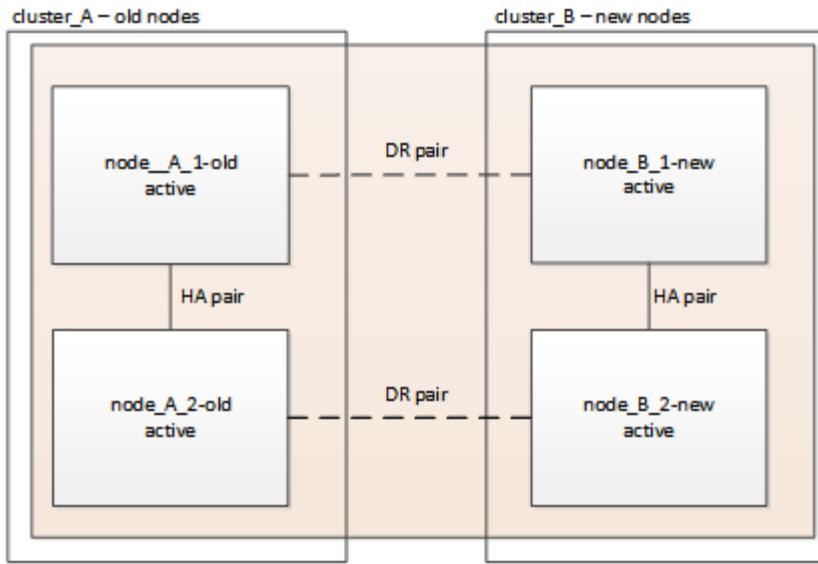
1. [Installing licenses for the new controller module](#)

## 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
• AFF A700	• AFF A900
• FAS9000	• 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 to 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.

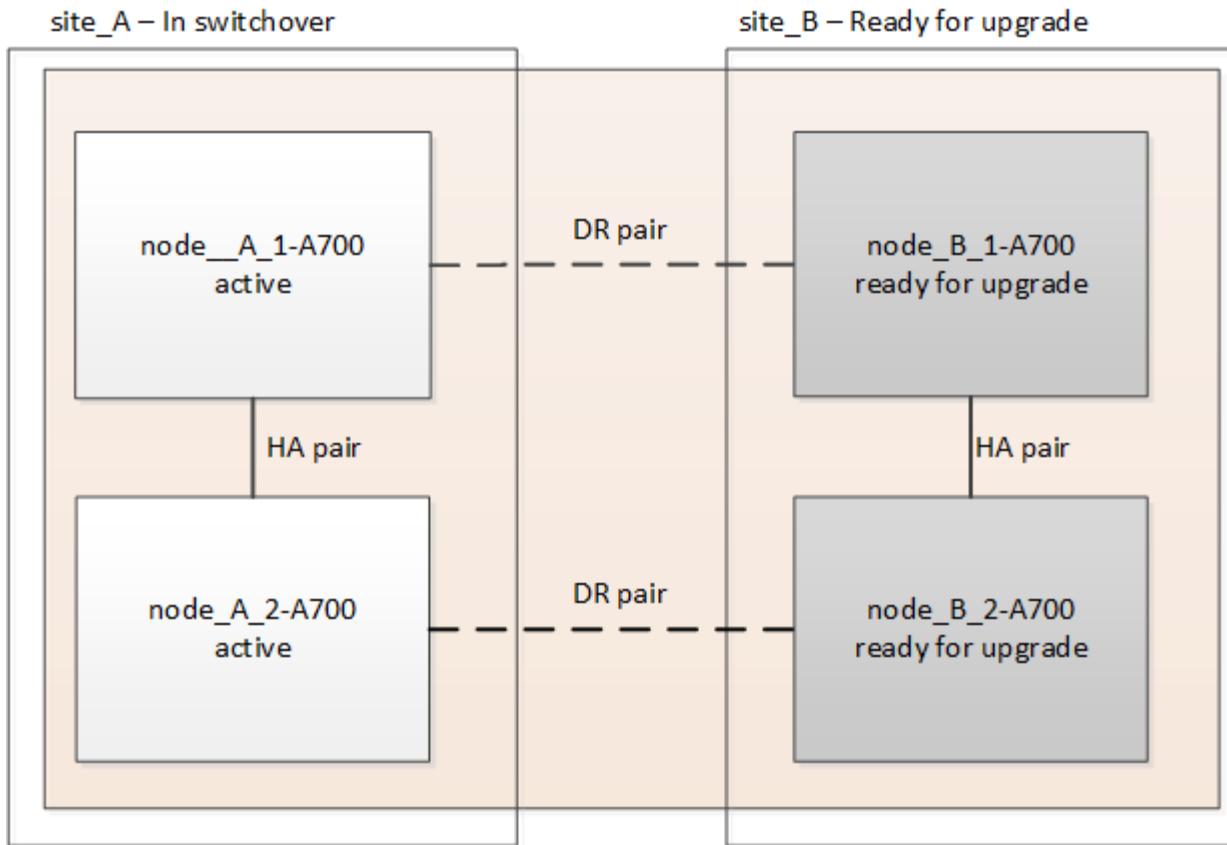
## Switch over the MetroCluster configuration

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

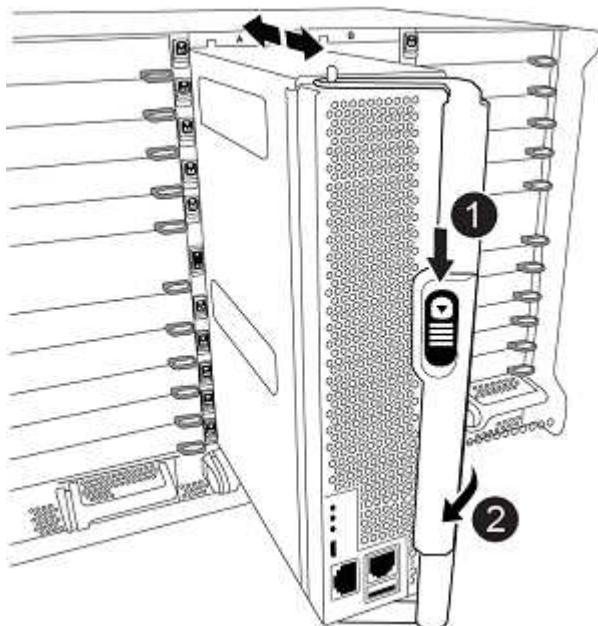
### Remove the AFF A700 or FAS9000 controller module

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.

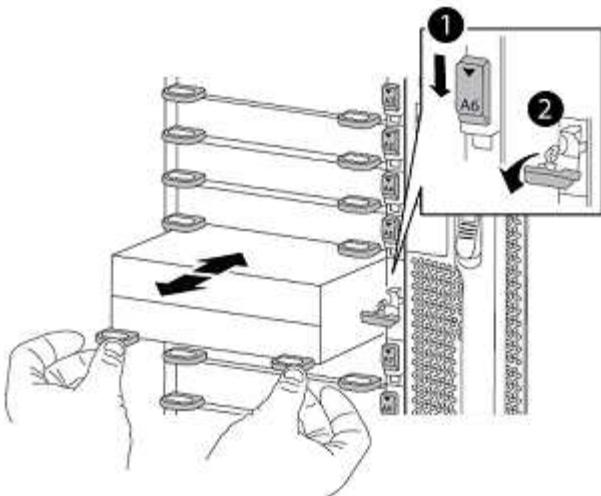
### Remove the AFF A700 or FAS9000 NVS module

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

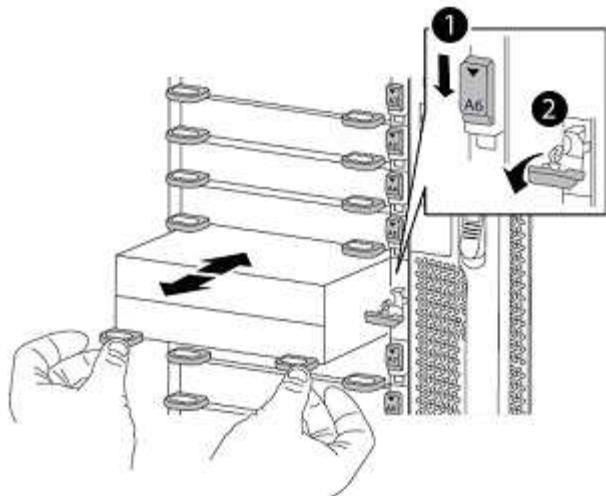
If you are not already grounded, properly ground yourself.

### 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.

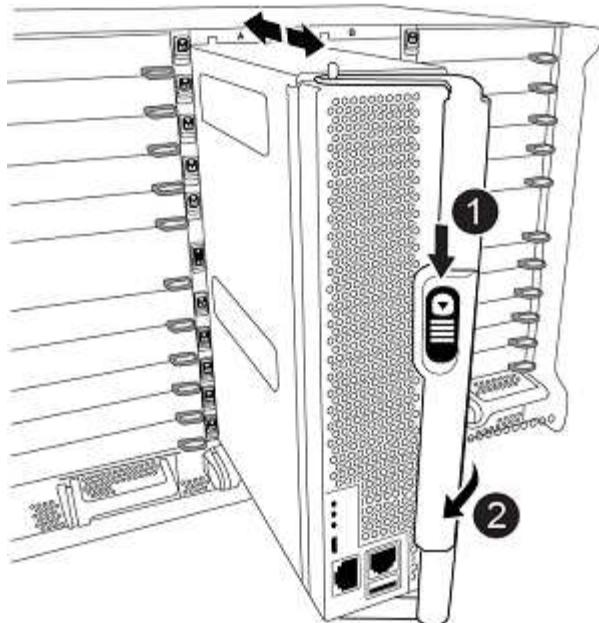
### Steps

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=&lt;filer_addr&gt; -mask=&lt;netmask&gt; -gw=&lt;gateway&gt; -dns=&lt;dns_addr&gt; domain=&lt;dns_domain&gt;</code>  <code>&lt;filer_addr&gt;</code> is the IP address of the storage system. <code>&lt;netmask&gt;</code> is the network mask of the storage system. <code>&lt;gateway&gt;</code> is the gateway for the storage system. <code>&lt;dns_addr&gt;</code> is the IP address of a name server on your network. This parameter is optional. <code>&lt;dns_domain&gt;</code> 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 <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>

### 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	<code>mcc-2n</code>

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>	ha-config modify controller mcc-2n ha-config modify chassis mcc-2n
<b>Four or eight nodes</b>	ha-config modify controller mcc ha-config modify chassis mcc

## 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)  Pool11  PZHYN0MD
node_B_1-A700(4068741254)  node_B_1-A700(4068741254)
rr18:9.126L49 node_B_1-A700(4068741254)  Pool11  PPG3J5HA
node_B_1-A700(4068741254)  node_B_1-A700(4068741254)
rr18:8.126L21 node_B_1-A700(4068741254)  Pool11  PZHTDSZD
node_B_1-A700(4068741254)  node_B_1-A700(4068741254)
rr18:8.126L2  node_B_1-A700(4068741254)  Pool0  S0M1J2CF
node_B_1-A700(4068741254)  node_B_1-A700(4068741254)
rr18:8.126L3  node_B_1-A700(4068741254)  Pool0  S0M0CQM5
node_B_1-A700(4068741254)  node_B_1-A700(4068741254)
rr18:9.126L27 node_B_1-A700(4068741254)  Pool0  S0M1PSDW
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 reassigned -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  S0M1J2CF
node_B_1-A900(1574774970)  node_B_1-A900(1574774970)
rr18:9.126L29 node_B_1-A900(1574774970)  Pool0  S0M0CQM5
node_B_1-A900(1574774970)  node_B_1-A900(1574774970)
rr18:8.126L1 node_B_1-A900(1574774970)  Pool0  S0M1PSDW
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_xxxxxxxxxx_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:

- View the broadcast domains:

```
network port broadcast-domain show
```

- Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

- Add the physical port that will host the intercluster LIFs to the corresponding Broadcast domain.
- Modify intercluster LIFs to use the new physical port as home port.
- 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 vserver-name
```

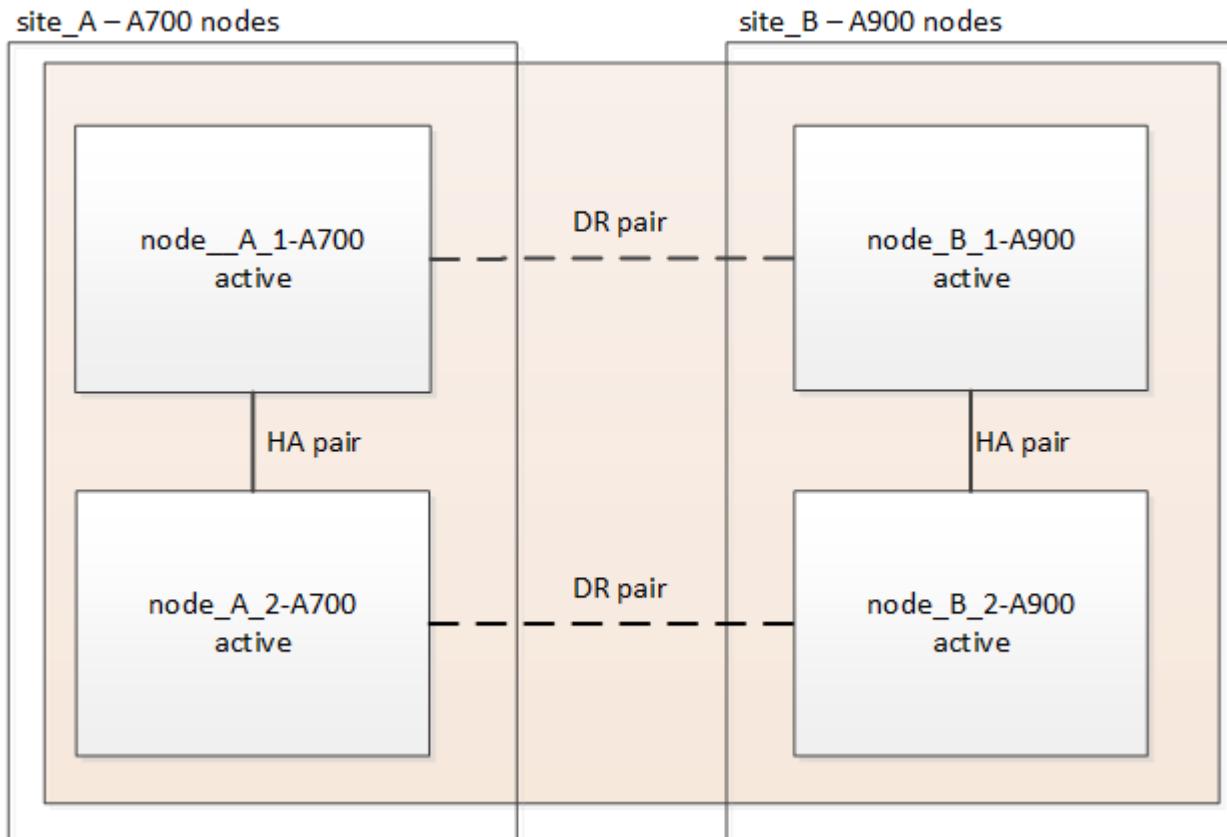
Perform this step on all SVMs as required.

## 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.
  - Verify that the new nodes are represented correctly.
  - 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      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.
```

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:

- fcp adapter show -instance
- fcp 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 data1if1 -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 a1a -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 node-name:portname, node-name: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: tar -zvxf ontap-version_image.tgz</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](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>

## 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)  Pool0  S0M1J2CF
node_B_1-old(4068741254)  node_B_1-old(4068741254)
rr18:8.126L3  node_B_1-old(4068741254)  Pool0  S0M0CQM5
node_B_1-old(4068741254)  node_B_1-old(4068741254)
rr18:9.126L27 node_B_1-old(4068741254)  Pool0  S0M1PSDW
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  S0M1J2CF
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:9.126L29  node_B_1-new(1574774970)  Pool0  S0M0CQM5
node_B_1-new(1574774970)  node_B_1-new(1574774970)
rr18:8.126L1   node_B_1-new(1574774970)  Pool0  S0M1PSDW
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”  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. 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, 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).

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:

- View the broadcast domains:

```
network port broadcast-domain show
```

- Add any ports to a broadcast domain as needed.

[Add or remove ports from a broadcast domain](#)

- Add the physical port that will host the intercluster LIFs to the corresponding broadcast domain.
- Modify intercluster LIFs to use the new physical port as home port.
- 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](#)

- 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](#)

- 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>

<p>External key management</p>	<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="#">Restore external key management encryption keys</a>.</p>
--------------------------------	--

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](#).

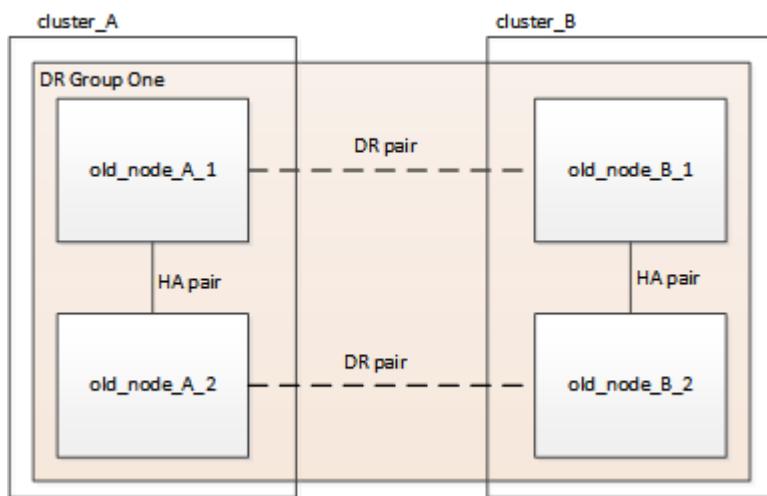
## Perform the refresh procedure

Use the following steps to refresh the MetroCluster FC configuration.

### Steps

1. Gather information from the old nodes.

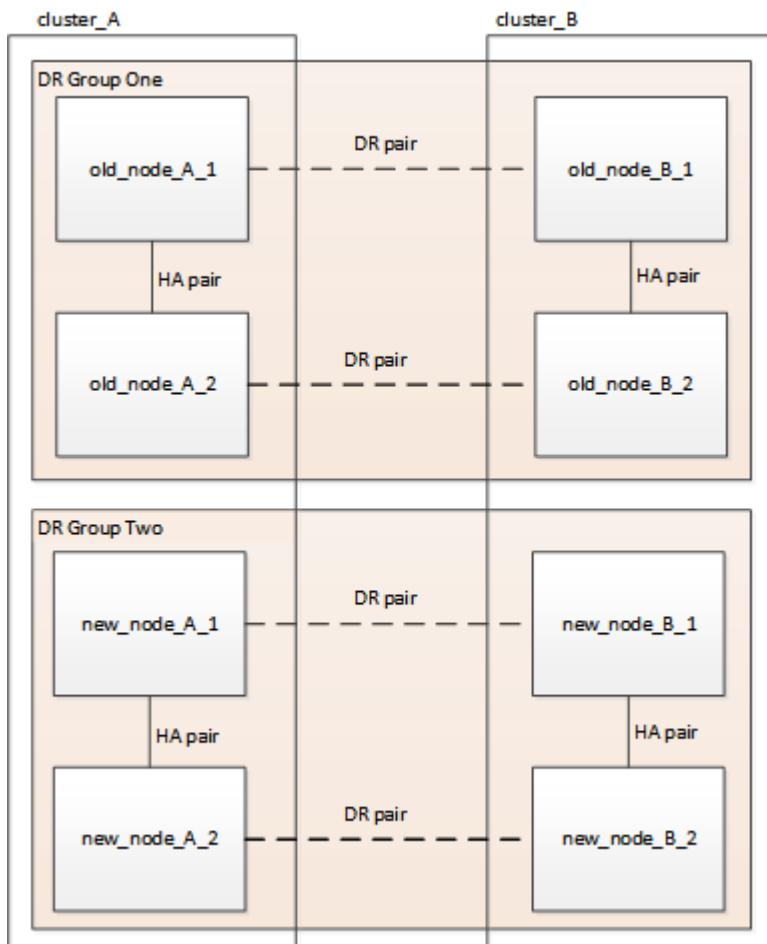
At this stage, the four-node configuration appears as shown in the following image:



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](#).

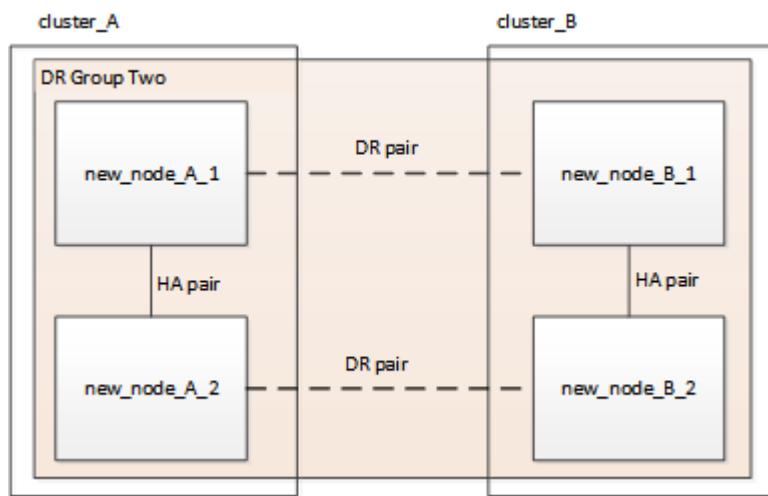
 You might choose to mirror the aggregate when or after it is created.

- 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.

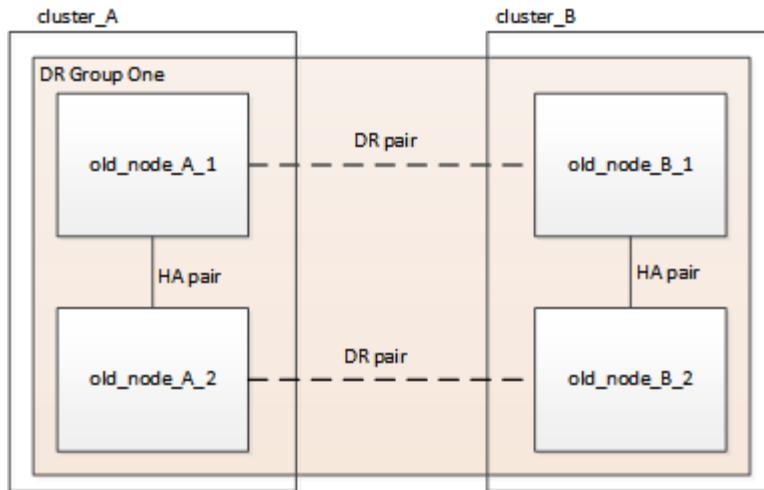
### Steps

1. Verify that you have a default broadcast domain created on the old nodes.

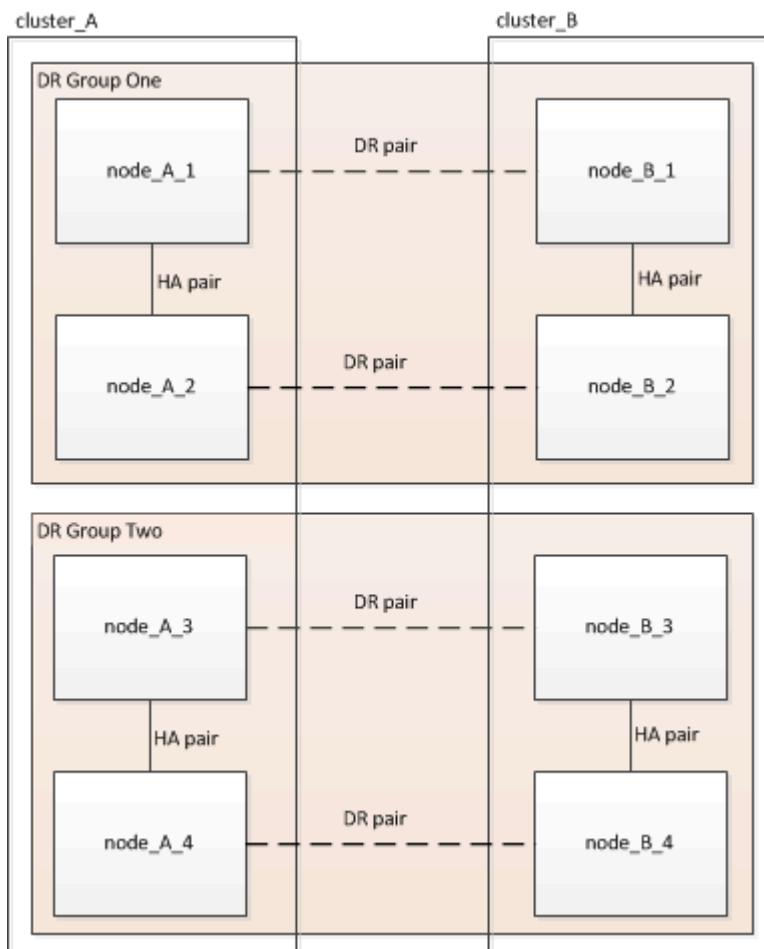
When you add new nodes to an existing cluster without a default broadcast domain, node management LIFs are created for the new nodes using universal unique identifiers (UUIDs) instead of the expected names. For more information, see the Knowledge Base article [Node management LIFs on newly-added nodes 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 -monitor -name cluster_A</code> command to remove the MetroCluster configuration. In the following example, “cluster_A” is removed from the software:<pre>NetApp MetroCluster Tiebreaker :&gt; monitor remove -monitor -name cluster_A Successfully removed monitor from NetApp MetroCluster Tiebreaker software.</pre></li><li>2. Confirm that the MetroCluster configuration is removed correctly by using the Tiebreaker CLI <code>monitor show -status</code> command.<pre>NetApp MetroCluster Tiebreaker :&gt; monitor show -status</pre></li></ol>
Mediator	Issue the following command from the ONTAP prompt:  <code>metrocluster configuration-settings mediator remove</code>
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:

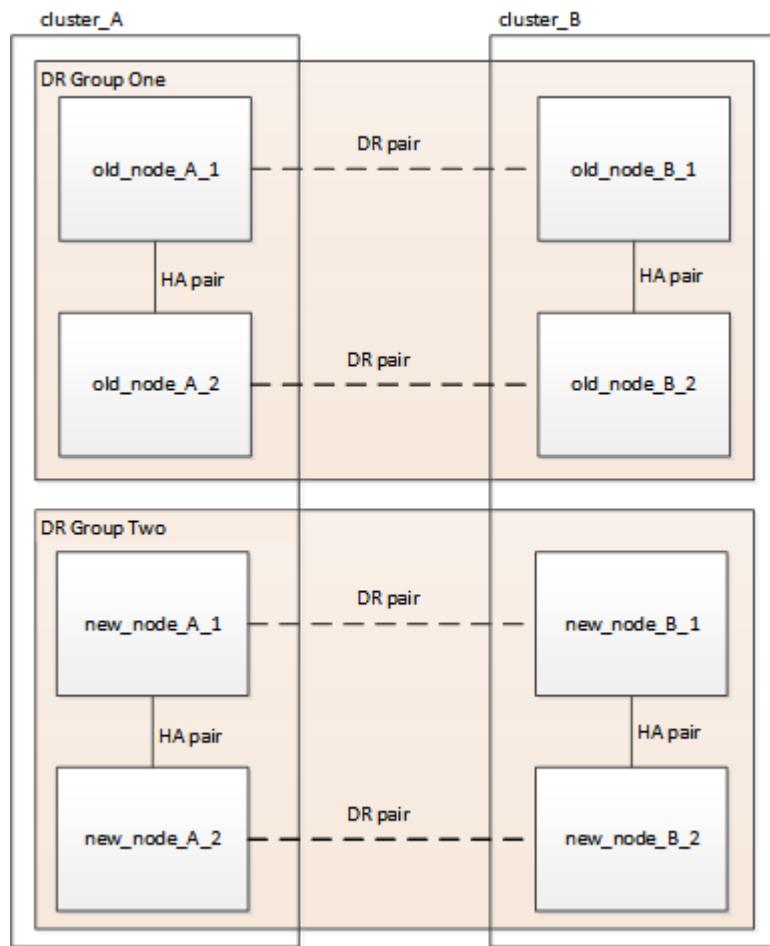
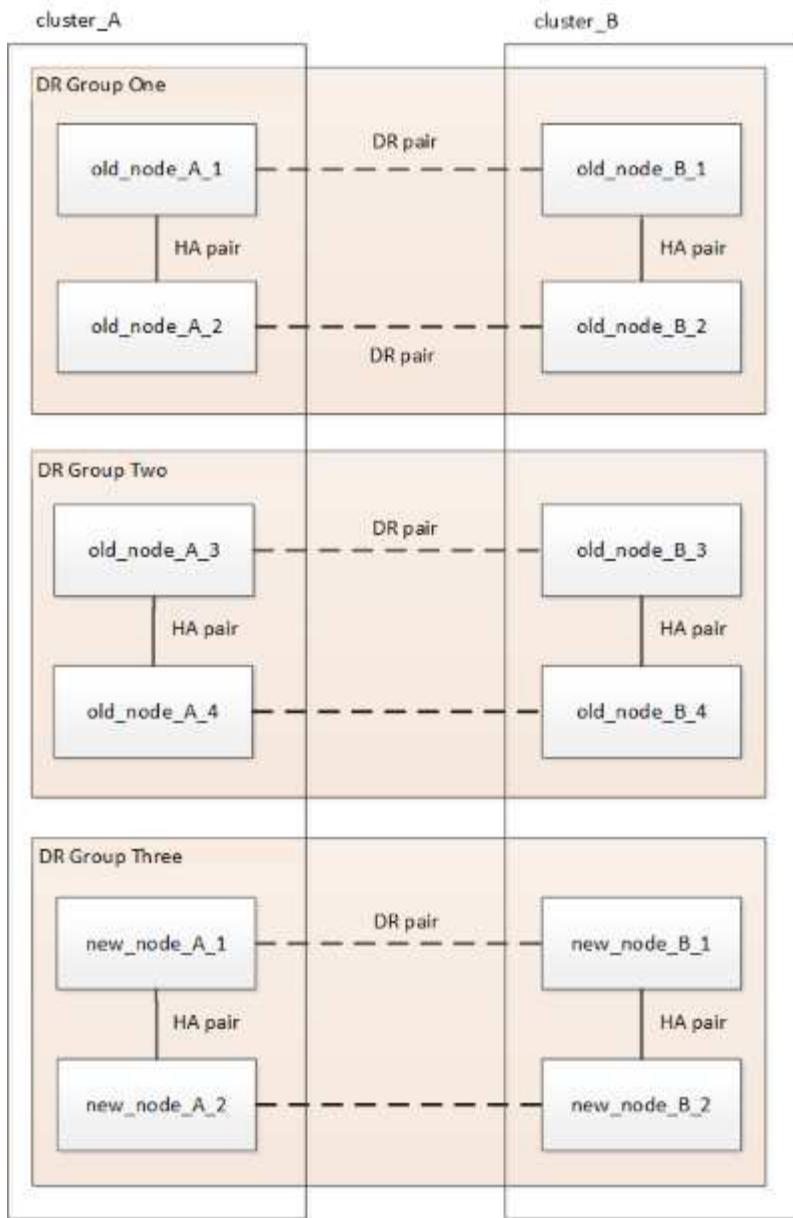


Figure 1. Temporary eight-node configuration



**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		Takeover	
Node	Partner	Possible	State Description
<hr/>			
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:

- 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.

- 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:

- 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
```

- Modify the cluster\_A peer IP address:

```
cluster peer modify -cluster cluster_A -peer-addrs node_A_3_IP -address
-family ipv4
```

- 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
```

- Modify the cluster\_B peer IP address:

```
cluster peer modify -cluster cluster_B -peer-addrs node_B_3_IP -address
-family ipv4
```

- Verify that the cluster peer IP address is updated for each cluster:

- 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

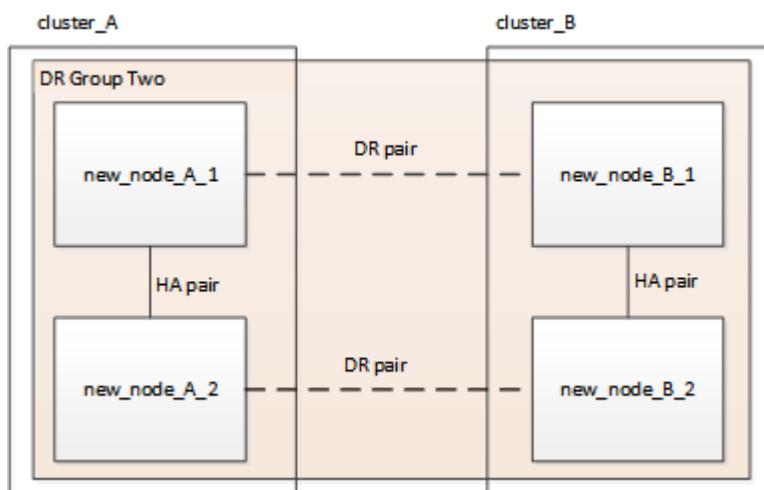
        Peer Cluster Name: cluster_A
        Remote Intercluster Addresses: 172.21.178.188,
172.21.178.196 <<<<<< Should reflect the modified address
        Availability of the Remote Cluster: Available
        Remote Cluster Name: cluster_A
        Active IP Addresses: 172.21.178.196,
172.21.178.188
        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::>

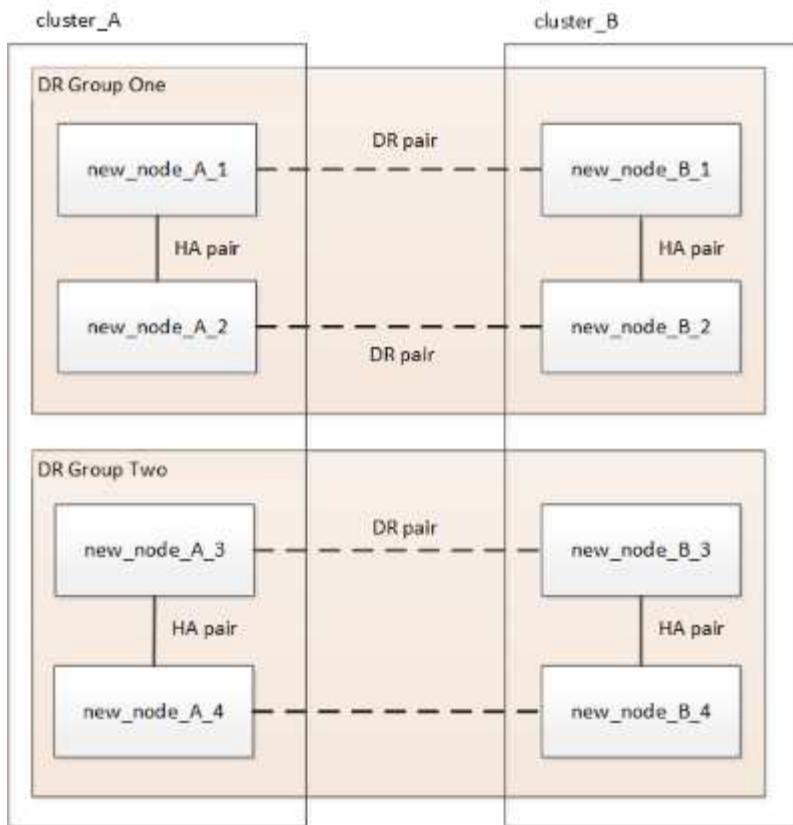
```

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	Configure ONTAP Mediator from a MetroCluster IP configuration 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.

- Issue the following command:

```
system node autosupport invoke -node * -type all -message MAINT=end
```

- 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

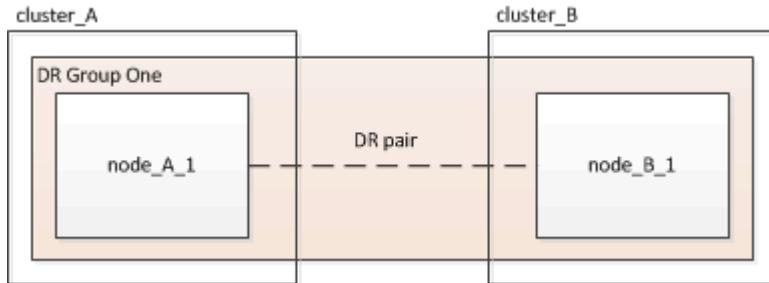
- You must have available FC switch ports to accommodate the new controllers and any new bridges.
- Verify that you have a default broadcast domain created on the old nodes.

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](#).

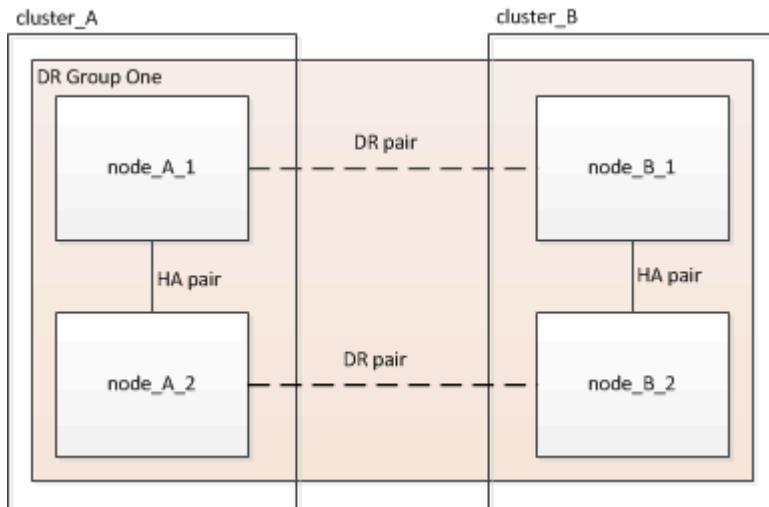
- You need the admin password and access to an FTP or SCP server.

## 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:



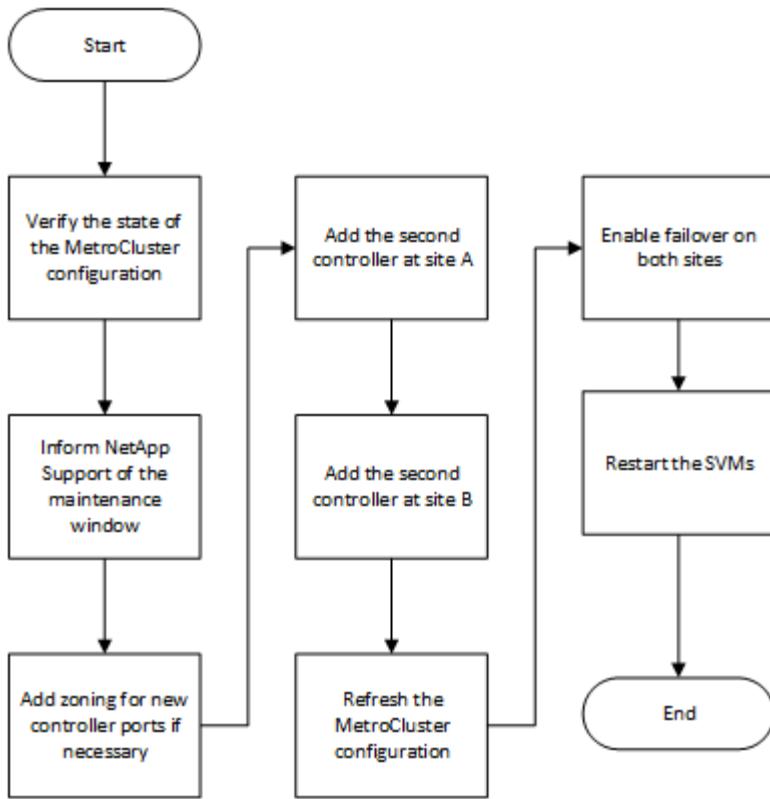
- Both sites must be expanded equally.

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          Size     Available Used%  State   #Vols  Nodes
RAID Status
-----
-----  -----
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).

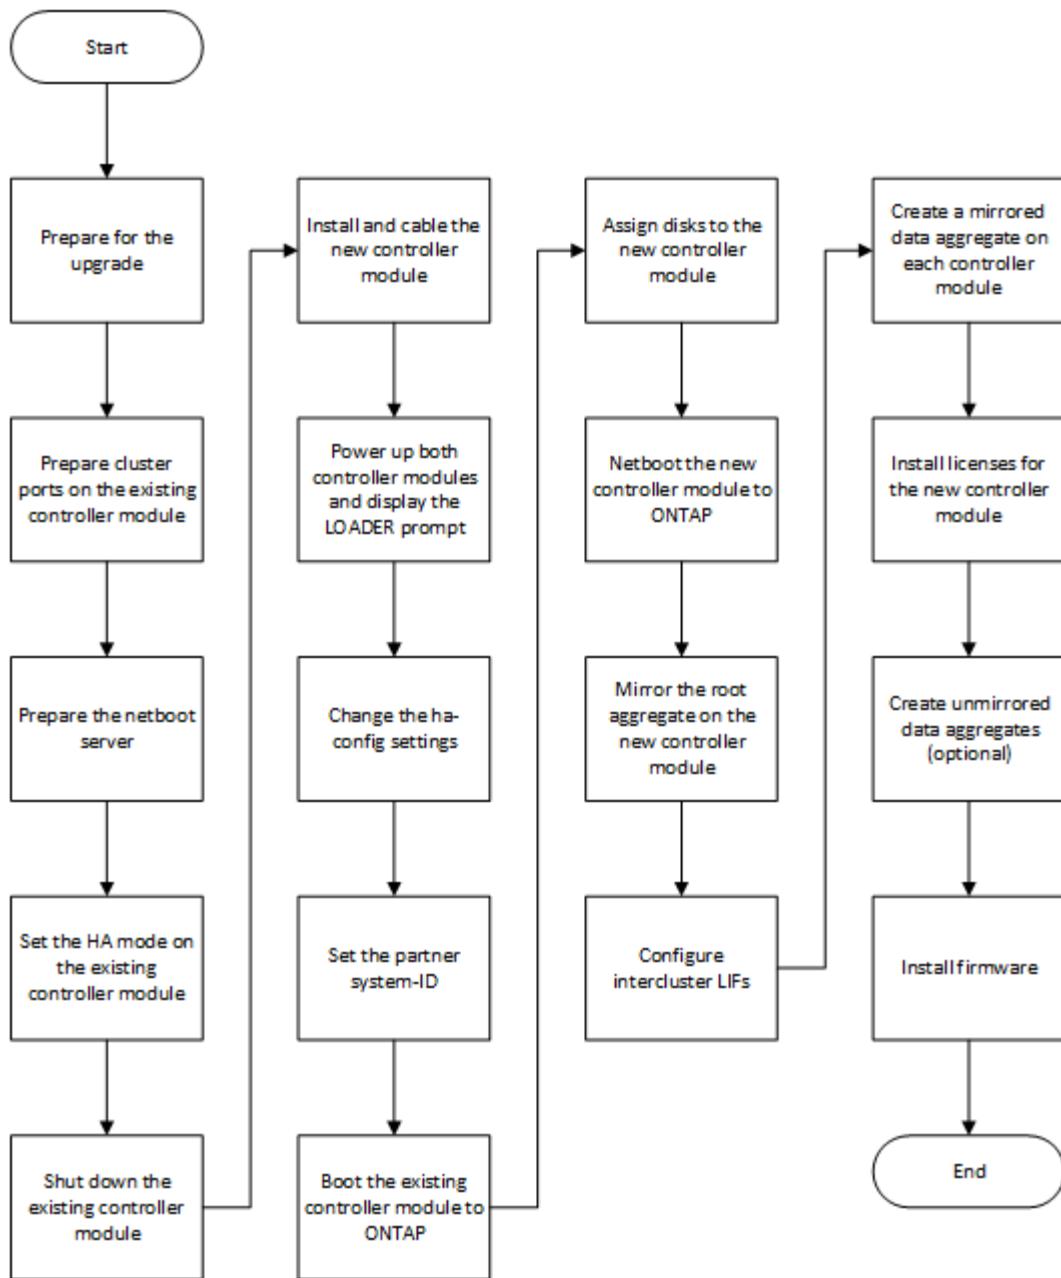


This configuration is not supported on all systems.

- 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_ctrlr>`.

- 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_ctrlr>`.
- 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 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 a1a -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 -ipspace Cluster -broadcast-domain Cluster -ports
port-id, port-id, port-id...
```

For example:

```
broadcast-domain add-ports -ipspace 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 ela -auto true
```

9. If you are creating a two-node switchless cluster, enable the switchless cluster networking mode:

- 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 (`*>`).

- Enable the switchless cluster networking mode:

```
network options switchless-cluster modify -enabled true
```

- 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...
FAS2200, FAS2500, FAS3200, FAS6200, FAS/AFF8000 series systems	<p>Extract the contents of the &lt;ontap_version&gt;_image.tgz 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>
All other systems	<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>

3. Determine the IP address of the existing controller module.

This address is referred to later in this procedure as *ip-address-of-existing controller*.

4. 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

1. 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 NVMMEM.

## 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...
To an empty bay to create a single-chassis HA pair and the system belongs to one of the following platforms:  • FAS8200 • 80xx	<ol style="list-style-type: none"><li>Remove the blank plate in the rear of the chassis that covers the empty bay that will contain the new controller module.</li><li>Gently push the controller module halfway into the chassis.  To prevent the controller module from automatically booting, do not fully seat it in the chassis until later in this procedure.</li></ol>
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.	Install the new system in the rack or system cabinet.

3. Cable the cluster network connections, as necessary:

- 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](#)

- 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.

- 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:

<b>If the controller modules are...</b>	<b>Then...</b>
---	----------------

In the same chassis	<ol style="list-style-type: none"> <li>Verify that the new controller module is <b>not</b> fully inserted into the bay.</li> </ol> <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> <ol style="list-style-type: none"> <li>Connect the power and turn on the power supplies so that the existing controller module receives power.</li> <li>Interrupt the boot process on the existing controller module by pressing Ctrl-C.</li> <li>Push the new controller module firmly into the bay.</li> </ol> <p>When fully seated, the new controller module receives power and automatically boots.</p> <ol style="list-style-type: none"> <li>Interrupt the boot process by pressing Ctrl-C.</li> <li>Tighten the thumbscrew on the cam handle, if present.</li> <li>Install the cable management device, if present.</li> <li>Bind the cables to the cable management device with the hook and loop strap.</li> </ol>
In separate chassis	<ol style="list-style-type: none"> <li>Turn on the power supplies on the existing controller module.</li> <li>Interrupt the boot process by pressing Ctrl-C.</li> <li>Repeat these steps for the new controller module</li> </ol>

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

- Change the ha-config setting of the existing controller module:
  - 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	<code>ifconfig e0M -auto</code>
Not available	<code>ifconfig e0M -addr=filer_addr -mask=netmask -gw=gateway -dns=dns_addr -domain=dns_domain</code>  <i>filer_addr</i> is the IP address of the storage system.  <i>netmask</i> is the network mask of the storage system.  <i>gateway</i> is the gateway for the storage system.  <i>dns_addr</i> is the IP address of a name server on your network.  <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.   Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.

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
                                         Speed
                                         (Mbps)
Node  Port      IPspace      Broadcast Domain Link   MTU    Admin/Oper
-----  -----  -----  -----  -----  -----  -----
-----  -----
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
          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
```

## 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

- Display a list of available spares:

```
storage disk show -spare -owner <node_name>
```

- 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
                                         Speed (Mbps)
Node  Port      IPspace    Broadcast Domain Link      MTU    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-ctlr	new-ctlr	true	Connected to new-ctlr
new-ctlr	old-ctlr	true	Connected to old-ctlr
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.

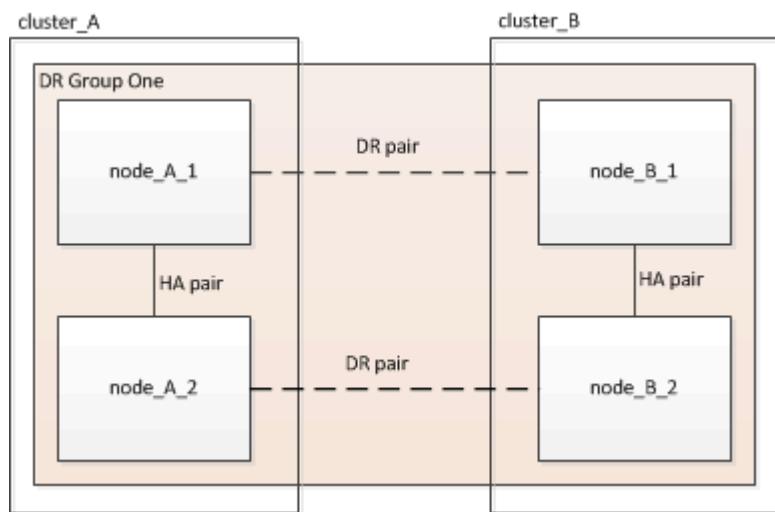
### About this task

- The nodes must be running ONTAP 9 in a MetroCluster FC configuration.

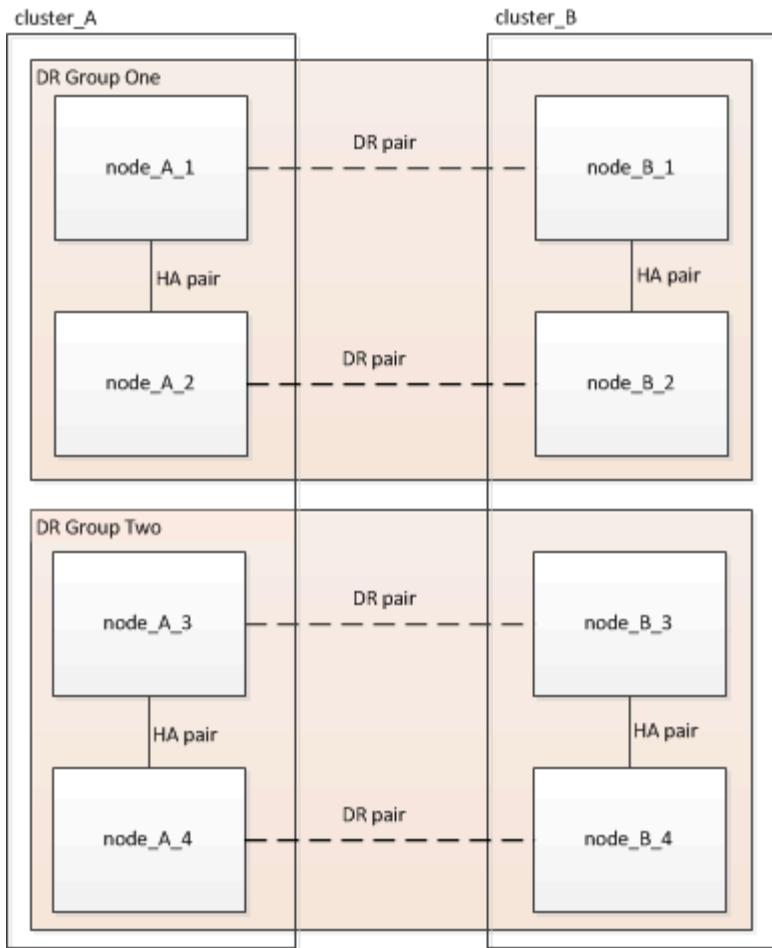
This procedure is not supported on earlier versions of ONTAP or in MetroCluster IP configurations.

- 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.

**i**

- 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:

- 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 (\*>).

- 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-name shelf-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
                                         Speed
                                         (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
          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
          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
                                         Speed (Mbps)
Node  Port      IPspace   Broadcast Domain Link    MTU     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 into 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 bridge-name</code>
9.4 and earlier	<code>storage bridge add -address bridge-ip-address -name bridge-name</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:

```
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

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 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 AFF A220	FAS500f AFF C250 AFF A250	FAS50	AFF C30 AFF A30	FAS8200 AFF A300	AFF A320	FAS8300 AFF C400 AFF A400	FAS8700
Eight-node 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 AFF C400 AFF A400										
	FAS8700										
	AFF C60										
	AFF A50										
	FAS70										
	FAS9000 AFF A700										
	AFF A70										
	AFF A70										
	FAS800 AFF A800										
	FAS9500 AFF A900										
	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 AFF A700	AFF A70	AFF C800 AFF A800	FAS9500 AFF A900	AFF C80	FAS90 AFF A90	AFF A1K
Eight-node 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										
	AFF A400										
	FAS8700										
	AFF C60										
	AFF A50										
	FAS70				Note 1						
	FAS9000										
	AFF A700										
	AFF A70										
	AFF C800										
	AFF A800										
	FAS9500										
	AFF A900										
	AFF C80										
	FAS90										
	AFF A90										
	AFF A1K										

## 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 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

          Speed (Mbps)  Health
Port      IPspace      Broadcast Domain Link MTU Admin/Oper Status
-----  -----
e0a       Cluster       Cluster        up   9000  auto/10000 healthy
e0b       Cluster       Cluster        up   9000  auto/10000 healthy

Node: node_A_2-old

          Speed (Mbps)  Health
Port      IPspace      Broadcast Domain Link MTU 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 have a Status Admin/Oper of up/up

```

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-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::>

```

## 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	<p>Issue the following command from the ONTAP prompt:</p> <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

```

```

        Home Port: ela
        172.17.28.11    255.255.255.0    -
completed

        Home Port: e1b
        172.17.29.11    255.255.255.0    -
completed

cluster_B

node_B_3-new
        Home Port: ela
        172.17.28.13    255.255.255.0    -
completed

        Home Port: e1b
        172.17.29.13    255.255.255.0    -
completed

node_B_3-new
        Home Port: ela
        172.17.28.12    255.255.255.0    -
completed

        Home Port: e1b
        172.17.29.12    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
----- ----- ----- ----- ----- ----- -----
```

DR	Source	Destination						
Group	Cluster	Node	Network	Address	Network	Address	Partner	Type
Config	State							
-----	-----	-----	-----	-----	-----	-----	-----	-----
1	cluster_B							

	node_B_2-old	Home Port: e1a	172.17.28.13	172.17.28.12	HA Partner
completed		Home Port: e1a	172.17.28.13	172.17.28.11	DR Partner
completed		Home Port: e1a	172.17.28.13	172.17.28.10	DR Auxiliary
completed		Home Port: e1b	172.17.29.13	172.17.29.12	HA Partner
completed		Home Port: e1b	172.17.29.13	172.17.29.11	DR Partner
completed		Home Port: e1b	172.17.29.13	172.17.29.10	DR Auxiliary
completed	node_B_1-old	Home Port: e1a	172.17.28.12	172.17.28.13	HA Partner
completed		Home Port: e1a	172.17.28.12	172.17.28.10	DR Partner
completed		Home Port: e1a	172.17.28.12	172.17.28.11	DR Auxiliary
completed		Home Port: e1b	172.17.29.12	172.17.29.13	HA Partner
completed		Home Port: e1b	172.17.29.12	172.17.29.10	DR Partner
completed		Home Port: e1b	172.17.29.12	172.17.29.11	DR Auxiliary
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	Source	Destination
	Group Cluster Node	Network Address	Network Address	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_rhal_a1
node_a_1
2.20.3          546.9GB  20   3 SAS      aggregate  aggr0_rhal_a2
node_a_2
2.20.5          546.9GB  20   5 SAS      aggregate  rhal_a1_aggr1
node_a_1
2.20.6          546.9GB  20   6 SAS      aggregate  rhal_a1_aggr1
node_a_1
2.20.7          546.9GB  20   7 SAS      aggregate  rhal_a2_aggr1
node_a_2
2.20.10         546.9GB  20  10 SAS      aggregate  rhal_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)

                                         Usable

Physical
    Position    Disk
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-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:  <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:  <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		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

Physical		Position	Disk	Type	Size
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                               Configuration  DR
Group Cluster Node             State        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.
```

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](#).

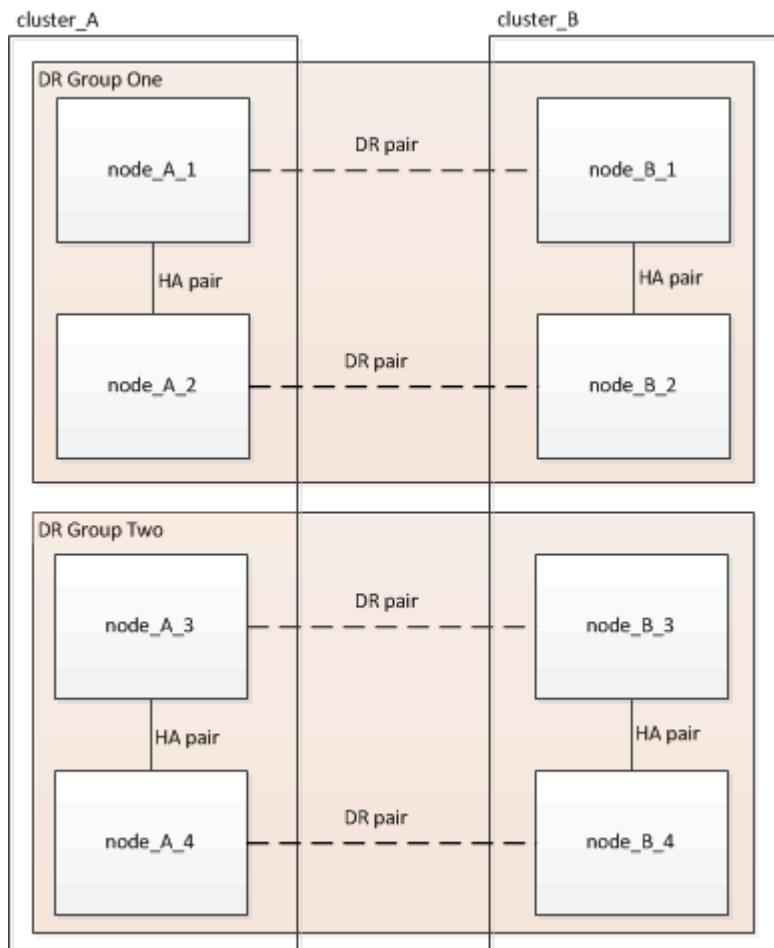
- Enable session logging for any CLI session. For instructions on how to enable session logging, review the "Logging Session Output" section in the Knowledge Base article [How to configure PuTTY for optimal connectivity to ONTAP systems](#).

## Remove the DR group nodes from each cluster

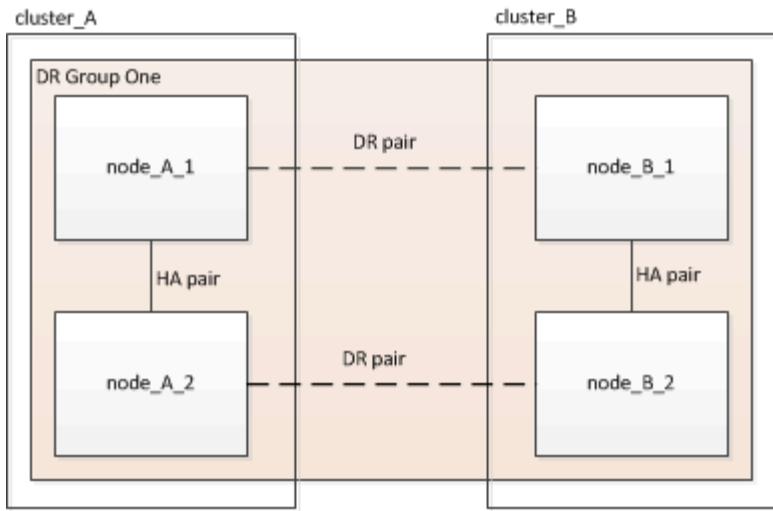
This procedure is supported beginning with ONTAP 9.8. For systems running ONTAP 9.7 or earlier, see the Knowledge Base article: [How to remove a DR group from a MetroCluster configuration](#).

### 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>
      
```
- g. 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>
<a href="#">Fabric-attached MetroCluster installation and configuration</a>	<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>

Stretch MetroCluster installation and configuration	<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>
MetroCluster management and disaster recovery	<ul style="list-style-type: none"> <li>• Understanding the MetroCluster configuration</li> <li>• Switchover, healing and switchback</li> <li>• Disaster recovery</li> </ul>
Maintain MetroCluster Components	<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 configuration.</li> </ul>
MetroCluster Upgrade, Transition, and Expansion	<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>
MetroCluster Tiebreaker Software installation and configuration	<ul style="list-style-type: none"> <li>• Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software</li> </ul>
ONTAP Hardware Systems Documentation	<ul style="list-style-type: none"> <li>• Hot-adding a disk shelf</li> <li>• Hot-removing a disk shelf</li> </ul>



The standard storage shelf maintenance procedures can be used with MetroCluster IP configurations.

Copy-based transition	<ul style="list-style-type: none"><li>• Transitioning data from 7-Mode storage systems to clustered storage systems</li></ul>
ONTAP concepts	<ul style="list-style-type: none"><li>• How mirrored aggregates work</li></ul>

## 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.