



Upgrade your controllers

ONTAP MetroCluster

NetApp
February 06, 2026

Table of Contents

Upgrade your controllers	1
Switch over the MetroCluster IP configuration	1
Remove interface configurations and uninstall the old MetroCluster IP controllers	2
Set up the new MetroCluster IP controllers	6
Restore the HBA configuration and set the HA state of the MetroCluster IP controller and chassis	9
Restore the HBA configuration	9
Set the HA state on the new controllers and chassis	10
Remove internal drives from the chassis on the new controller	11
Update the switch RCFs and set the MetroCluster IP bootarg values	12
Update the switch RCFs to accommodate the new platforms	12
Set the MetroCluster IP bootarg variables	12
Reassign the root aggregate disks to the new MetroCluster IP controller module	19
Boot the new MetroCluster IP controllers and restore LIF configuration	22
Boot the new controllers	22
Verify and restore LIF configuration	25
Switch back the MetroCluster IP configuration	26

Upgrade your controllers

Switch over the MetroCluster IP configuration

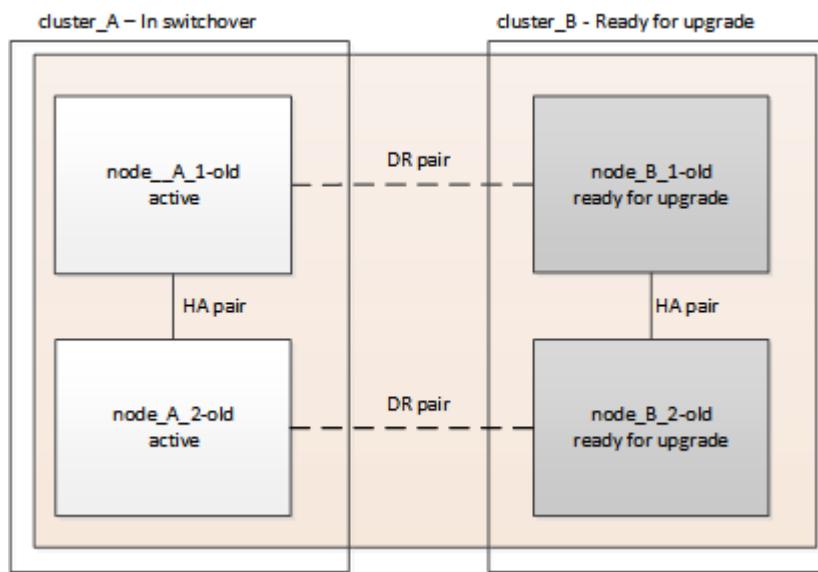
You switch over the configuration to site_A so that the platforms on site_B can be upgraded.

About this task

This task must be performed on site_A.

After you complete this task:

- cluster_A is active and serving data for both sites.
- cluster_B is inactive and ready to begin the upgrade process.



Steps

1. Switch over the MetroCluster configuration to site_A so that site_B's nodes can be upgraded:

a. Issue the following command on cluster_A:

```
metrocluster switchover -controller-replacement true
```

The operation can take several minutes to complete.

b. Monitor the switchover operation:

```
metrocluster operation show
```

c. After the operation is complete, confirm that the nodes are in switchover state:

```
metrocluster show
```

d. Check the status of the MetroCluster nodes:

```
metrocluster node show
```

Automatic healing of aggregates after negotiated switchover is disabled during a controller upgrade.

What's next?

[Remove interface configurations and uninstall the old controllers.](#)

Remove interface configurations and uninstall the old MetroCluster IP controllers

Verify the correct LIF placement. Then remove the VLANs and interface groups on the old controllers and physically uninstall the controllers.

About this task

- You perform these steps on the old controllers (node_B_1-old, node_B_2-old).
- You require the information you gathered in [Map ports from the old nodes to the new nodes](#) for use in this procedure.

Steps

1. Boot the old nodes and log in to the nodes:

```
boot_ontap
```

2. If the system you are upgrading to uses **shared cluster/HA ports**, verify that the MetroCluster IP interfaces are using supported IP addresses.

Use the following information to determine whether the new system uses shared cluster/HA ports:

Shared cluster/HA ports

The systems listed in the following table use shared cluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A20• AFF A30• AFF C30• AFF A50• AFF C60• AFF C80• AFF A70• AFF A90• AFF A1K	<ul style="list-style-type: none">• FAS50• FAS70• FAS90

Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A150, ASA A150• AFF A220• AFF C250, ASA C250• AFF A250, ASA A250• AFF A300• AFF A320• AFF C400, ASA C400• AFF A400, ASA A400• AFF A700• AFF C800, ASA C800• AFF A800, ASA A800• AFF A900, ASA A900	<ul style="list-style-type: none">• FAS2750• FAS500f• FAS8200• FAS8300• FAS8700• FAS9000• FAS9500

a. Verify the IP addresses of the MetroCluster interfaces on the old controllers:

```
metrocluster configuration-settings interface show
```

b. If the MetroCluster interfaces are using 169.254.17.x or 169.254.18.x IP addresses, refer to the [Knowledge Base article "How to modify the properties of a MetroCluster IP interface"](#) to modify the interface IP addresses before you proceed with the upgrade.



Upgrading to any system that uses **shared cluster/HA ports** isn't supported if the MetroCluster interfaces are configured with 169.254.17.x or 169.254.18.x IP addresses.

3. Modify the intercluster LIFs on the old controllers to use a different home port than the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers.



This step is required for a successful upgrade.

The intercluster LIFs on the old controllers must use a different home port than the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers. For example, when you upgrade to AFF A90 controllers, the HA interconnect ports are e1a and e7a, and the MetroCluster IP DR interconnect ports are e2b and e3b. You must move the intercluster LIFs on the old controllers if they are hosted on ports e1a, e7a, e2b, or e3b.

For port distribution and allocation on the new nodes, refer to the [Hardware Universe](#).

- a. On the old controllers, view the intercluster LIFs:

```
network interface show -role intercluster
```

Take one of the following actions depending on whether the intercluster LIFs on the old controllers use the same ports as the ports used for HA interconnect or MetroCluster IP DR interconnect on the new controllers.

If the intercluster LIFs...	Go to...
Use the same home port	Substep b
Use a different home port	Step 4

- b. Modify the intercluster LIFs to use a different home port:

```
network interface modify -vserver <vserver> -lif <intercluster_lif> -home -port <port-not-used-for-ha-interconnect-or-mcc-ip-dr-interconnect-on-new-nodes>
```

- c. Verify that all intercluster LIFs are on their new home ports:

```
network interface show -role intercluster -is-home false
```

The command output should be empty, indicating that all intercluster LIFs are on their respective home ports.

- d. Revert any LIFs that aren't on their home ports:

```
network interface revert -lif <intercluster_lif>
```

Repeat the command for each intercluster LIF that isn't on the home port.

4. Assign the home port of all data LIFs on the old controller to a common port that is the same on both the old and new controller modules.



If the old and new controllers don't have a common port, you don't need to modify the data LIFs. Skip this step and go directly to [Step 5](#).

a. Display the LIFs:

```
network interface show
```

All data LIFs including SAN and NAS are admin up and operationally down because those are up at switchover site (cluster_A).

b. Review the output to find a common physical network port that is the same on both the old and new controllers that is not used as a cluster port.

For example, e0d is a physical port on old controllers and is also present on new controllers. e0d is not used as a cluster port or otherwise on the new controllers.

For port usage for platform models, see the [Hardware Universe](#)

c. Modify all data LIFs to use the common port as the home port:

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-port  
<port-id>
```

In the following example, this is "e0d".

For example:

```
network interface modify -vserver vs0 -lif 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.

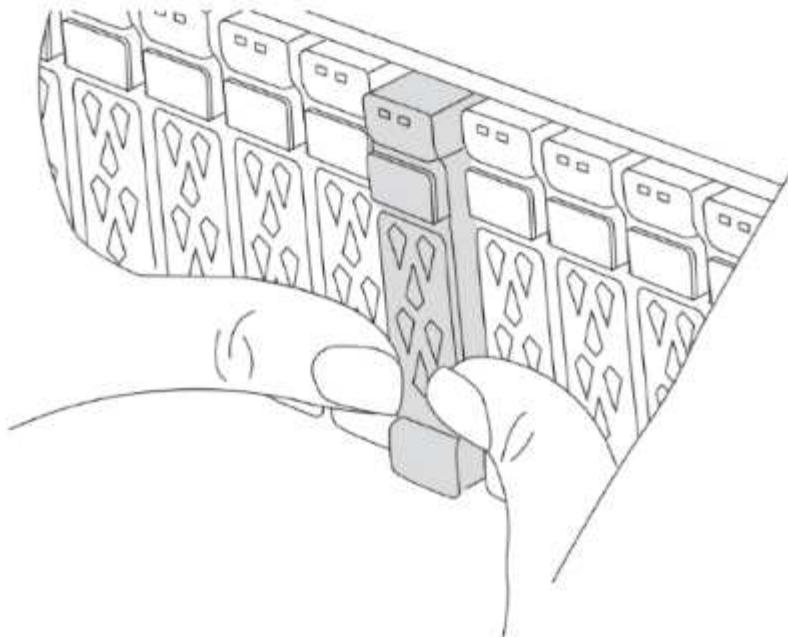
Steps

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">• An AFF A150 to an AFF A20 system• An AFF A220 to an AFF A20 system	Controller to external shelf conversion
<ul style="list-style-type: none">• An AFF A800 to an AFF A90 system• An AFF C800 to an AFF C80 system• An AFF A250 to an AFF A30 system• An AFF C250 to an AFF C30 system• An AFF A70 to an AFF A90 system	Controller module replacement
Any other controller upgrade combinations	All other upgrades

Controller to external shelf conversion

If your original MetroCluster IP controllers are AFF A150 or AFF A220 models, you can convert the AFF A150 or AFF A220 HA pair to a DS224C drive shelf and then attach it to the new nodes.

For example, when upgrading from an AFF A150 or AFF A220 system to an AFF A20 system, you can convert the AFF A150 or AFF A220 HA pair to a DS224C shelf by swapping the AFF A150 or AFF A220 controller modules with IOM12 modules.

Steps

1. Replace the controller modules in the node you are converting with IOM12 shelf modules.

[Hardware Universe](#)

2. Set the drive shelf ID.

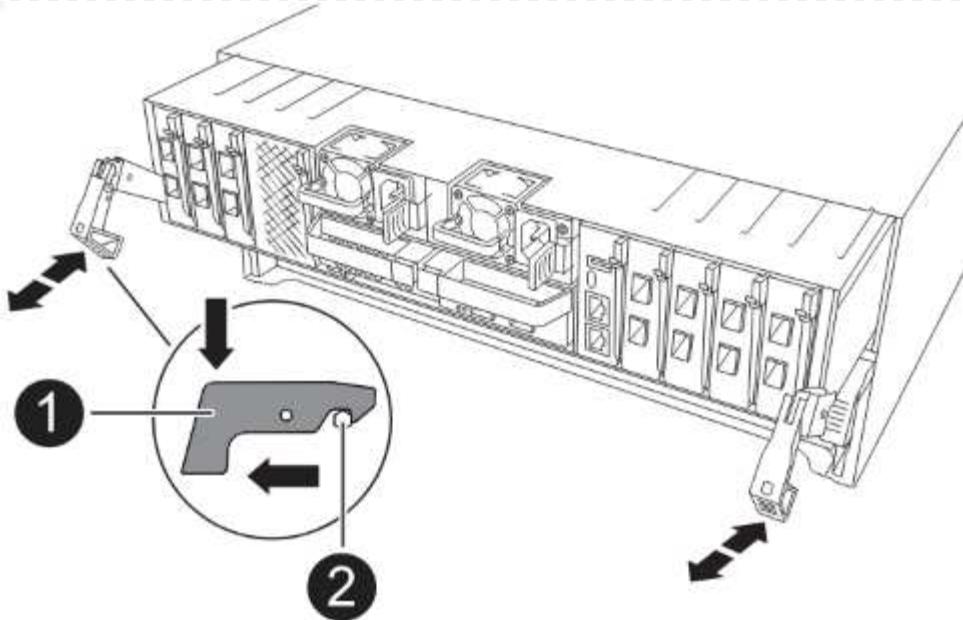
Each drive shelf, including the chassis, requires a unique ID.

3. Reset other drive shelf IDs as needed.
4. Power off the shelves.
5. Cable the converted drive shelf to a SAS port on the new system, and, if you are using out-of-band ACP cabling, to the ACP port on the new node.
6. Turn on the power to the converted drive shelf and any other drive shelves attached to the new nodes.
7. Turn on the power to the new nodes, and then interrupt the boot process on each node by pressing Ctrl-C to access the boot environment prompt.

Controller module replacement

Installing the new controllers separately is not applicable for upgrades of integrated systems with disks and controllers in the same chassis, for example, from an AFF A800 system to an AFF A90 system. You must swap the new controller modules and I/O cards after powering off the old controllers, as shown in the image below.

The following example image is for representation only, the controller modules and I/O cards can vary between systems.



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.

If you have this type of HBA and desired mode...	Use this command...
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">• AFF A20• AFF A30• AFF C30• AFF A50• AFF C60• AFF C80• AFF A70• AFF A90• AFF A1K	<ul style="list-style-type: none">• FAS50• FAS70• FAS90

Shared MetroCluster/HA ports

The systems listed in the following table use shared MetroCluster/HA ports:

AFF and ASA systems	FAS systems
<ul style="list-style-type: none">• AFF A150, ASA A150• AFF A220• AFF C250, ASA C250• AFF A250, ASA A250• AFF A300• AFF A320• AFF C400, ASA C400• AFF A400, ASA A400• AFF A700• AFF C800, ASA C800• AFF A800, ASA A800• AFF A900, ASA A900	<ul style="list-style-type: none">• FAS2750• FAS500f• FAS8200• FAS8300• FAS8700• FAS9000• FAS9500

Steps

1. At the `LOADER>` prompt, set the following bootargs on the new nodes at site_B:

The steps you follow depend on the ports used by the new platform model.

Systems that use shared cluster/HA ports

- a. Set the following bootargs:

```
setenv bootarg.mcc.port_a_ip_config <local-IP-address/local-IP-  
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id>
```

```
setenv bootarg.mcc.port_b_ip_config <local-IP-address/local-IP-  
mask,0,0,DR-partner-IP-address,DR-aux-partnerIP-address,vlan-id>
```



If the interfaces are using a default VLAN ID, the `vlan-id` parameter is not required.

The following example sets the values for `node_B_1-new` using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12,130
```

The following example sets the values for `node_B_2-new` using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,0,172.17.26.12,172.17.26.13,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,0,172.17.27.12,172.17.27.13,130
```

The following example sets the values for `node_B_1-new` using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,0,172.17.26.13,172.17.26.12  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,0,172.17.27.13,172.17.27.12
```

The following example sets the values for `node_B_2-new` using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,0,172.17.26.12,172.17.26.13  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,0,172.17.27.12,172.17.27.13
```

Systems that use shared MetroCluster/HA ports

- a. Set the following bootargs:

```
setenv bootarg.mcc.port_a_ip_config <local-IP-address/local-IP-  
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-  
address,vlan-id>
```

```
setenv bootarg.mcc.port_b_ip_config <local-IP-address/local-IP-  
mask,0,HA-partner-IP-address,DR-partner-IP-address,DR-aux-partnerIP-  
address,vlan-id>
```



If the interfaces are using a default VLAN ID, the `vlan-id` parameter is not required.

The following example sets the values for `node_B_1-new` using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12,130
```

The following example sets the values for `node_B_2-new` using VLAN 120 for the first network and VLAN 130 for the second network:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13,120  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13,130
```

The following example sets the values for `node_B_1-new` using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.10/23,0,172.17.26.11,172.17.26.13,172.17.26.12  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.10/23,0,172.17.27.11,172.17.27.13,172.17.27.12
```

The following example sets the values for `node_B_2-new` using default VLANs for all MetroCluster IP DR connections:

```
setenv bootarg.mcc.port_a_ip_config  
172.17.26.11/23,0,172.17.26.10,172.17.26.12,172.17.26.13  
setenv bootarg.mcc.port_b_ip_config  
172.17.27.11/23,0,172.17.27.10,172.17.27.12,172.17.27.13
```

2. At the new nodes' LOADER prompt, set the UUIDs:

```
setenv bootarg.mgwd.partner_cluster_uuid <partner-cluster-UUID>  
setenv bootarg.mgwd.cluster_uuid <local-cluster-UUID>  
setenv bootarg.mcc.pri_partner_uuid <DR-partner-node-UUID>  
setenv bootarg.mcc.aux_partner_uuid <DR-aux-partner-node-UUID>  
setenv bootarg.mcc_iscsi.node_uuid <local-node-UUID>
```

a. Set the UUIDs on node_B_1-new:

The following example shows the commands for setting the UUIDs on node_B_1-new:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039  
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-  
00a098c9e55d  
setenv bootarg.mcc.pri_partner_uuid f37b240b-9ac1-11e7-9b42-  
00a098c9e55d  
setenv bootarg.mcc.aux_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-  
00a098ca379f  
setenv bootarg.mcc_iscsi.node_uuid f03cb63c-9a7e-11e7-b68b-  
00a098908039
```

b. Set the UUIDs on node_B_2-new:

The following example shows the commands for setting the UUIDs on node_B_2-new:

```
setenv bootarg.mgwd.cluster_uuid ee7db9d5-9a82-11e7-b68b-00a098908039  
setenv bootarg.mgwd.partner_cluster_uuid 07958819-9ac6-11e7-9b42-  
00a098c9e55d  
setenv bootarg.mcc.pri_partner_uuid bf8e3f8f-9ac4-11e7-bd4e-  
00a098ca379f  
setenv bootarg.mcc.aux_partner_uuid f37b240b-9ac1-11e7-9b42-  
00a098c9e55d  
setenv bootarg.mcc_iscsi.node_uuid aa9a7a7a-9a81-11e7-a4e9-  
00a098908c35
```

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 Restore onboard key management encryption keys .
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 Restore external key management encryption keys .

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

- a. View the broadcast domains:

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

- b. If a new broadcast domain is created for the data ports on the newly upgraded controllers, delete the broadcast domain:



Only delete the new broadcast domain. Don't delete any of the broadcast domains that existed before starting the upgrade.

```
broadcast-domain delete -broadcast-domain <broadcast_domain_name>
```

- c. Add ports to a broadcast domain as needed.

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

- d. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might differ from the old node.

[Create a VLAN](#)

[Combine physical ports to create interface groups](#)

Verify and restore LIF configuration

Verify that LIFs are hosted on appropriate nodes and ports as mapped out at the beginning of the upgrade procedure.

About this task

- This task is performed on site_B.
- See the port mapping plan you created in [Map ports from the old nodes to the new nodes](#).



You must verify that the location of the data LIFs is correct on the new nodes before you perform a switchback. When you switchback the configuration, ONTAP attempts to resume traffic on the home port used by the LIFs. I/O failure can occur when the home port connection to the switch port and VLAN is incorrect.

Steps

1. Verify that LIFs are hosted on the appropriate node and ports before switchback.

- a. Change to the advanced privilege level:

```
set -privilege advanced
```

- b. Display the LIFs, and confirm that each data LIF is using the correct home port:

```
network interface show
```

- c. Modify any LIFs that aren't using the correct home port:

```
network interface modify -vserver <svm-name> -lif <data-lif> -home-port
```

```
<port-id>
```

If the command returns an error, you can override the port configuration:

```
vserver config override -command "network interface modify -vserver <svm-name> -home-port <active_port_after_upgrade> -lif <lif_name> -home-node <new_node_name>"
```

When entering the network interface modify command within the vserver config override command, you cannot use the tab autocomplete feature. You can create the network interface modify using autocomplete and then enclose it in the vserver config override command.

- d. Confirm that all data LIFs are now on the correct home port:

```
network interface show
```

- e. Return to the admin privilege level:

```
set -privilege admin
```

2. Revert the interfaces to their home node:

```
network interface revert * -vserver <svm-name>
```

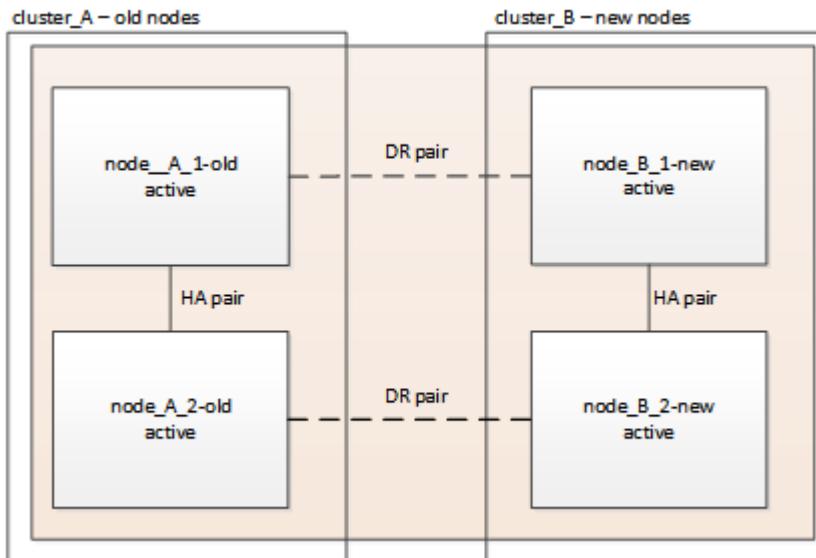
Perform this step on all SVMs as required.

What's next?

[Switchback the MetroCluster configuration.](#)

Switch back the MetroCluster IP configuration

Perform the switchback operation to return the MetroCluster configuration to normal operation. The nodes on site_A are still awaiting upgrade.



Steps

1. Issue the `metrocluster node show` command on site_B and check the output.
 - a. Verify that the new nodes are represented correctly.
 - b. Verify that the new nodes are in "Waiting for switchback state."
2. Perform the healing and switchback by running the required commands from any node in the active cluster (the cluster that is not undergoing upgrade).
 - a. Heal the data aggregates:
`metrocluster heal aggregates`
 - b. Heal the root aggregates:
`metrocluster heal root`
 - c. Switchback the cluster:
`metrocluster switchback`
3. Check the progress of the switchback operation:
`metrocluster show`

The switchback operation is still in progress when the output displays `waiting-for-switchback`:

```
cluster_B::> metrocluster show
Cluster          Entry Name      State
-----
Local: cluster_B  Configuration state configured
                  Mode           switchover
                  AUSO Failure Domain -
Remote: cluster_A Configuration state configured
                  Mode           waiting-for-switchback
                  AUSO Failure Domain -
```

The switchback operation is complete when the output displays `normal`:

```
cluster_B::> metrocluster show
Cluster          Entry Name      State
-----
Local: cluster_B  Configuration state configured
                  Mode           normal
                  AUSO Failure Domain -
Remote: cluster_A Configuration state configured
                  Mode           normal
                  AUSO Failure Domain -
```

If a switchback takes a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command. This command is at the

advanced privilege level.

What's next?

[Complete the upgrade.](#)

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—with 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.