



Upgrade or expand the MetroCluster configuration



ONTAP

Fourth Edition (March 2024)

© Copyright Lenovo 2021, 2024.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant to a General Services Administration (GSA) contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Contents

1. Start here - Choose your procedure	1
1.1. Start here: Choose between controller upgrade, system refresh, or expansion	1
1.2. Choose a controller upgrade procedure	2
1.2.1. Choosing a procedure that uses the switchover and switchback process	2
Supported MetroCluster IP controller upgrades	2
Supported MetroCluster FC controller upgrades	3
1.3. Choosing a system refresh method	4
Supported MetroCluster FC tech refresh combinations	4
Supported MetroCluster IP tech refresh combinations	4
1.4. Choose an expansion procedure	6
2. Upgrading controllers in a MetroCluster FC configuration using switchover and switchback	7
2.1. Supported platform combinations	7
2.2. About this task	7
2.3. Preparing for the upgrade	8
2.3.1. Verifying the health of the MetroCluster configuration	8
2.3.2. Mapping ports from the old nodes to the new nodes	10
2.3.3. Gathering information before the upgrade	11
2.3.4. Removing the existing configuration from the Tiebreaker or other monitoring software	13
2.3.5. Sending a custom AutoSupport message prior to maintenance	13
2.4. Switching over the MetroCluster configuration	14
2.5. Preparing the network configuration of the old controllers	16
2.6. Removing the old platforms	18
2.7. Configuring the new controllers	18
2.7.1. Setting up the new controllers	18
2.7.2. Netbooting the new controllers	19
2.7.3. Clearing the configuration on a controller module	20
2.7.4. Restoring the HBA configuration	20
2.7.5. Setting the HA state on the new controllers and chassis	21
2.7.6. Reassigning root aggregate disks	22
2.7.7. Booting up the new controllers	24
2.7.8. Verifying LIF configuration	28
2.7.9. Install the new licenses	29
2.8. Switching back the MetroCluster configuration	29
2.9. Checking the health of the MetroCluster configuration	31
2.10. Upgrading the nodes on cluster_A	32
2.11. Sending a custom AutoSupport message after maintenance	32
2.12. Restoring Tiebreaker monitoring	32
3. Upgrading controllers in a four-node MetroCluster FC configuration using switchover and switchback with "system controller replace" commands (ONTAP 9.10.1 and later)	33

3.1. Supported platform combinations	33
3.2. About this task	33
3.3. Preparing for the upgrade	34
3.3.1. Gathering information before the upgrade	35
3.3.2. Removing the existing configuration from the Tiebreaker or other monitoring software	37
3.4. Replacing the old controllers and booting up the new controllers	38
3.4.1. Preparing the network configuration of the old controllers	38
3.4.2. Setting up the new controllers	40
3.4.3. Netbooting the new controllers	40
3.4.4. Clearing the configuration on a controller module	41
3.4.5. Restoring the HBA configuration	42
3.4.6. Reassigning root aggregate disks	43
3.4.7. Booting up the new controllers	45
3.5. Completing the upgrade	48
3.5.1. Restoring Tiebreaker monitoring	49
4. Upgrading controllers in a MetroCluster IP configuration using switchover and switchback (ONTAP 9.8 and later)	50
4.1. Platforms supported by this procedure	50
4.2. About this task	50
4.3. Workflow for upgrading controllers in an MetroCluster IP configuration	51
4.4. Preparing for the upgrade	52
4.4.1. Updating the MetroCluster switch RCF files before upgrading controllers	52
4.4.2. Mapping ports from the old nodes to the new nodes	54
4.4.3. Netbooting the new controllers	55
4.4.4. Clearing the configuration on a controller module	57
4.4.5. Verifying MetroCluster health before site upgrade	57
4.4.6. Gathering information before the upgrade	58
4.4.7. Removing Mediator or Tiebreaker monitoring	61
4.4.8. Sending a custom AutoSupport message prior to maintenance	61
4.5. Switching over the MetroCluster configuration	62
4.6. Removing interface configurations and uninstalling the old controllers	63
4.7. Updating the switch RCFs to accommodate the new platforms	65
4.8. Configuring the new controllers	65
4.8.1. Setting up the new controllers	65
4.8.2. Restoring the HBA configuration	66
4.8.3. Setting the HA state on the new controllers and chassis	67
4.8.4. Setting the MetroCluster IP bootarg variables	67
4.8.5. Reassigning root aggregate disks	70
4.8.6. Booting up the new controllers	72
4.8.7. Verifying and restoring LIF configuration	75
4.9. Switching back the MetroCluster configuration	76

4.10. Checking the health of the MetroCluster configuration	77
4.11. Upgrading the nodes on cluster_A	78
4.12. Restoring Tiebreaker or Mediator monitoring	78
4.13. Sending a custom AutoSupport message after maintenance	79
5. Refreshing a four-node MetroCluster FC configuration	80
6. Refreshing a four-node or an eight-node MetroCluster IP configuration (ONTAP 9.8 and later)	83
7. Expand a four-node MetroCluster FC configuration to an eight-node configuration	92
7.1. Expanding a four-node MetroCluster FC configuration to an eight-node configuration	92
7.1.1. Supported platform combinations when adding a second DR group	94
7.2. Determining the new cabling layout	94
7.3. Racking the new equipment	95
7.4. Verifying the health of the MetroCluster configuration	95
7.5. Sending a custom AutoSupport message prior to adding nodes to the MetroCluster configuration	96
7.6. Recable and zone a switch fabric for the new nodes	97
7.6.1. Disconnecting the existing DR group from the fabric	97
7.6.2. Recable and reconfigure the switches	97
7.7. Configure ONTAP on the new controllers	99
7.7.1. Clearing the configuration on a controller module	99
7.7.2. Assigning disk ownership in AFA systems	99
7.7.3. Assigning disk ownership in non-AFA systems	101
7.7.4. Verifying the ha-config state of components	104
7.7.5. Booting the new controllers and joining them to the cluster	105
7.7.6. Configure the clusters into a MetroCluster configuration	107
Configure intercluster LIFs	107
8. Expanding a MetroCluster IP configuration	124
8.1. Example naming in this procedure	124
8.2. Supported platform combinations when adding a second DR group	124
8.3. Sending a custom AutoSupport message prior to maintenance	125
8.4. Verifying the health of the MetroCluster configuration	126
8.5. Removing the configuration from monitoring applications	128
8.6. Preparing the new controller modules	129
8.7. Upgrade RCF files	130
8.8. Joining the new nodes to the clusters	132
8.9. Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates	133
8.10. Finalizing the addition of the new nodes	142
9. Removing a Disaster Recovery group	148
9.1. Removing the DR group nodes from each cluster	149
10. Where to find additional information	153
10.1. MetroCluster and miscellaneous information	153
11. Appendix	155
11.1. Contacting support	155

11.2. Notices	155
11.3. Trademarks	156

Chapter 1. Start here - Choose your procedure

1.1. 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.

Upgrade type	Go to...
Controller upgrade	Choose a controller upgrade procedure
System refresh	Choose a system refresh procedure
Expansion	<ul style="list-style-type: none">• Four-node MetroCluster FC to eight• Four-node MetroCluster IP to eight

1.2. 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.

1.2.1. Choosing a procedure that uses the switchover and switchback process

Select your Current platform from the FC or IP table below. If the intersection of the Current platform row and Target platform column is blank, the upgrade is not supported.

Supported MetroCluster IP controller upgrades

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.

	Target MetroCluster IP platform				
		DM5000F DM5000H	DG5000 DM5100F	DM7000H DM7000F	DM7100H DM7100F
Source MetroCluster IP platform	DM5000F DM5000H				
	DG5000 DM5100F				
	DM7000H DM7000F				
	DM7100H DM7100F				

- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, the new controllers must be running the same ONTAP version as the old controllers.

Supported MetroCluster FC controller upgrades

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.

		Target MetroCluster FC platform			
		DM7000H	DM7000F	DM7100H	DM7100F
Source MetroCluster FC platform	DM7000H			Note 1	
	DM7000F				Note 1
	DM7100H				
	DM7100F				

- Note 1: Controller upgrades from DM7000F or DM700H platforms using onboard ports 0e and 0f as FC-VI connections are supported only on systems running ONTAP 9.9.1 or earlier.
- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, the new controllers must be running the same ONTAP version as the old controllers.

MetroCluster type	Upgrade method	ONTAP version	Procedure
IP	Upgrade with 'system controller replace' commands	9.13.1 and later	Link to procedure
FC	Upgrade with 'system controller replace' commands	9.10.1 and later	Link to procedure
FC	Manual upgrade with CLI commands	9.8 and later	Link to procedure
IP	Manual upgrade with CLI commands	9.8 and later	Link to procedure

1.3. 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 FC tech refresh combinations

		Target MetroCluster FC platform			
		DM7000H	DM7000F	DM7100H	DM7100F
Current MetroCluster FC platform	DM7000H				
	DM7000F				
	DM7100H				
	DM7100F				

- 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.
- 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 [Lenovo Press](#) for platform limits.

Supported MetroCluster IP tech refresh combinations

		Target MetroCluster IP platform			
		DM5000F DM5000H	DM5100F	DM7000H DM7000F	DM7100H DM7100F DG7000
Source MetroCluster IP platform	DM5000F				
	DM5000H				
	DM5100F				
	DM7000H				
	DM7000F				
	DM7100H DM7100F DG7000				

- 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.
- 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 [Lenovo Press](#) for platform limits.

Refresh method	Configuration type	ONTAP version	Procedure
<ul style="list-style-type: none"> Method: Expand the MetroCluster configuration and then remove the old nodes 	Four-node FC	9.6 and later	Link to procedure
<ul style="list-style-type: none"> Method: Expand the MetroCluster configuration and then remove the old nodes 	Four-node IP	9.8 and later	Link to procedure

1.4. 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 four-node MetroCluster FC to eight	Four-node FC	ONTAP 9.4 or later	Link to procedure
Method: Expand a four-node MetroCluster IP to eight	Four-node IP	ONTAP 9.9.1 and later	Link to procedure

Chapter 2. Upgrading 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.

2.1. 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.

2.2. 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:
 - 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.
- The licenses at both sites must match.
- This procedure applies to controller modules in a MetroCluster FC configuration (a 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.

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

2.3. Preparing 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.

2.3.1. Verifying the health of the MetroCluster configuration

You must verify the health and connectivity of the MetroCluster configuration prior to 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
```

### 2.3.2. Mapping 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 upgrade, you must recreate 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 [Lenovo Press](#).

Also identify the FC-VI card slot usage.

2. Plan your port usage and, if desired, fill in the following tables for reference for each of the new nodes.

You will refer to the table as you carry out the upgrade procedure.

| LIF       | node_A_1-old |          |                   | node_A_1-new |          |                   |
|-----------|--------------|----------|-------------------|--------------|----------|-------------------|
|           | Ports        | IPspaces | Broadcast domains | Ports        | IPspaces | Broadcast domains |
| Cluster 1 |              |          |                   |              |          |                   |
| Cluster 2 |              |          |                   |              |          |                   |
| Cluster 3 |              |          |                   |              |          |                   |



|                    |  |  |  |  |  |  |
|--------------------|--|--|--|--|--|--|
| Cluster 4          |  |  |  |  |  |  |
| Node management    |  |  |  |  |  |  |
| Cluster management |  |  |  |  |  |  |
| Data 1             |  |  |  |  |  |  |
| Data 2             |  |  |  |  |  |  |
| Data 3             |  |  |  |  |  |  |
| Data 4             |  |  |  |  |  |  |
| SAN                |  |  |  |  |  |  |
| Intercluster port  |  |  |  |  |  |  |

### 2.3.3. Gathering 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
1 Cluster_A Node_A_2-old 4068741260
4068741258
1 Cluster_B Node_B_1-old 4068741254
4068741256
1 Cluster_B Node_B_2-old 4068741256
4068741254
4 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:

- `fc adapter show -instance`
- `fc 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
```

### **2.3.4. Removing the existing configuration from the Tiebreaker or other monitoring software**

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration 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.

### **2.3.5. Sending a custom AutoSupport message prior to maintenance**

Before performing the maintenance, you should issue an AutoSupport message to notify Lenovo 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.

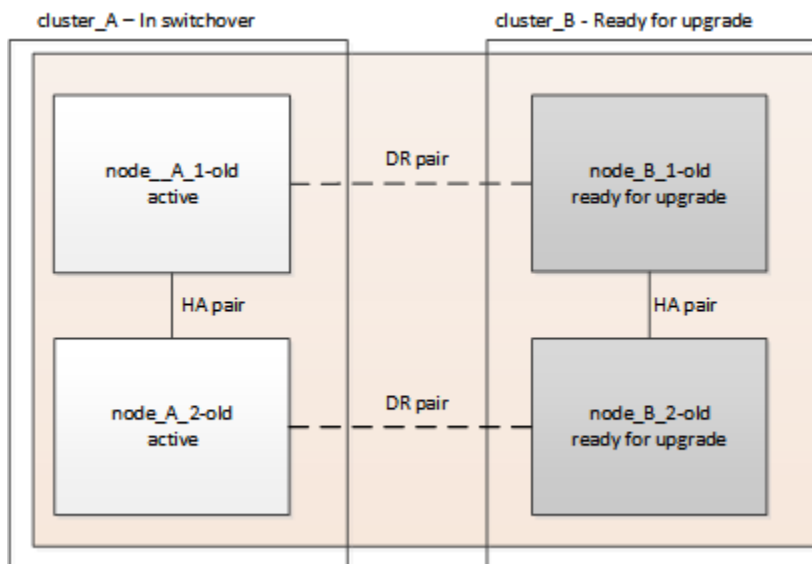
## 2.4. Switching over the MetroCluster configuration

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`

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: -
```

## 2.5. Preparing 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 [Lenovo Press](#)

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

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

In the following example, this is "e0d".

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

3. Modify broadcast domains to remove vlan and physical ports that need to be deleted:

```
broadcast-domain remove-ports -broadcast-domain broadcast-domain-name -ports node-name:port-id
```

Repeat this step for all VLAN and physical ports.

4. Remove any VLAN ports using cluster ports as member ports and ifgrps using cluster ports as member ports.

- a. Delete VLAN ports:

```
network port vlan delete -node node-name -vlan-name portid-vlandid
```

For example:

```
network port vlan delete -node node1 -vlan-name e1c-80
```

- b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node node-name -ifgrp interface-group-name -port portid
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp a1a -port e0d
```

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

```
network port broadcast-domain remove-ports -ip-space ip-space -broadcast-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..
```

- d. Modify interface group ports to use other physical ports as member as needed.:

```
ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id
```

5. Halt the nodes:

```
halt -inhibit-takeover true -node node-name
```

This step must be performed on both nodes.

## 2.6. Removing 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.

## 2.7. Configuring 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.

### 2.7.1. Setting 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.

[ThinkSystem DG and DM Series](#)

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 [Lenovo Press](#) 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.



6. Power up the new nodes and press Ctrl-C when prompted to display the LOADER prompt.

### 2.7.2. Netbooting 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.

1. Access [Lenovo Data Center Support](#) to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software 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.

| If the platform model is... | Then...                                                                                                                                                                                          |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| All systems                 | Your directory listing should contain a netboot folder with a kernel file: <code>ontap-version_image.tgz</code><br><br>You do not need to extract the <code>ontap-version_image.tgz</code> 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 by using this 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. Be sure to 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**
```

- 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}
// end include reference
```

### 2.7.3. Clearing the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the configuration.

- If necessary, halt the node to display the LOADER prompt: `halt`
- At the LOADER prompt, set the environmental variables to default values: `set-defaults`
- Save the environment: `saveenv`
- Restart the system by entering `bye`
- Press Ctrl+C to enter the LOADER prompt.
- At the LOADER prompt, launch the boot menu: `boot_ontap menu`
- At the boot menu prompt, type the following command to clear the configuration: `wipeconfig`

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

- At the boot menu, select option `5` to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

### 2.7.4. Restoring 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

- In Maintenance mode configure the settings for any HBAs in the system:
  - Check the current settings of the ports: `ucadmin show`
  - Update the port settings as needed.

| If you have this type of HBA and desired mode... | Use this command...                                         |
|--------------------------------------------------|-------------------------------------------------------------|
| CNA FC                                           | <code>ucadmin modify -m fc -t initiator 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> |

- Exit Maintenance mode:

`halt`

After you run the command, wait until the node stops at the LOADER prompt.

- Boot the node back into Maintenance mode to enable the configuration changes to take effect:

`boot_ontap maint`

- 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> |

### 2.7.5. Setting 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

- 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... |
|------------------------------------------|---------------------------|
| Four or eight nodes                      | <code>mcc</code>          |

- 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... |
|------------------------------------------|-------------------------|
|------------------------------------------|-------------------------|

|                     |                                                                         |
|---------------------|-------------------------------------------------------------------------|
| Four or eight nodes | <pre>ha-config modify controller mcc ha-config modify chassis mcc</pre> |
|---------------------|-------------------------------------------------------------------------|

### 2.7.6. Reassigning 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-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 HOME     | OWNER                    | POOL  | SERIAL NUMBER | HOME                     | DR |
|---------------|--------------------------|-------|---------------|--------------------------|----|
| rr18:9.126L44 | node_B_1-old(4068741254) | Poo11 | PZHYN0MD      | node_B_1-old(4068741254) |    |
| rr18:9.126L49 | node_B_1-old(4068741254) | Poo11 | PPG3J5HA      | node_B_1-old(4068741254) |    |
| rr18:8.126L21 | node_B_1-old(4068741254) | Poo11 | PZHTDSZD      | node_B_1-old(4068741254) |    |
| rr18:8.126L2  | node_B_1-old(4068741254) | Poo10 | S0M1J2CF      | node_B_1-old(4068741254) |    |
| rr18:8.126L3  | node_B_1-old(4068741254) | Poo10 | S0M0CQM5      | node_B_1-old(4068741254) |    |
| rr18:9.126L27 | node_B_1-old(4068741254) | Poo10 | S0M1PSDW      | node_B_1-old(4068741254) |    |

- 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
```

- 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) Poo11 PZHYN0MD node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:9.126L49 node_B_1-new(1574774970) Poo11 PPG3J5HA node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:8.126L21 node_B_1-new(1574774970) Poo11 PZHTDSZD node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:8.126L2 node_B_1-new(1574774970) Poo10 S0M1J2CF node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:9.126L29 node_B_1-new(1574774970) Poo10 S0M0CQM5 node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:8.126L1 node_B_1-new(1574774970) Poo10 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).

### 2.7.7. Booting 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"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option "11"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |

### ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, issue the boot menu command.

| If you are using...     | Issue this command at the boot menu prompt... |
|-------------------------|-----------------------------------------------|
| Onboard key management  | <code>recover_onboard_keymanager</code>       |
| External key management | <code>recover_external_keymanager</code>      |

5. If autoboot is enabled, interrupt autoboot by pressing CTRL-C.
6. From the boot menu, run option "6".



Option "6" will reboot the node twice before completing.

Respond “y” to the system id change prompts. Wait for the second reboot messages:

```
Successfully restored env file from boot media...
Rebooting to load the restored env file...
```

7. Double-check that the partner-sysid is correct:

```
printenv partner-sysid
```

If the partner-sysid is not correct, set it:

```
setenv partner-sysid partner-sysID
```

8. If root encryption is used, depending on the ONTAP version you are using, select the boot menu option or issue the boot menu command again for your key management configuration.



### ONTAP 9.8 and later

Beginning with ONTAP 9.8, select the boot menu option.

| If you are using...     | Select this boot menu option...                                                                                            |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option "10"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option "11"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |

Depending on the key manager setting, perform the recovery procedure by selecting option "10" or option "11", followed by option "6" at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option "1" (normal boot).

### ONTAP 9.7 and earlier

For ONTAP 9.7 and earlier, issue the boot menu command.

| If you are using...     | Issue this command at the boot menu prompt... |
|-------------------------|-----------------------------------------------|
| Onboard key management  | <code>recover_onboard_keymanager</code>       |
| External key management | <code>recover_external_keymanager</code>      |

You might need to issue the `recover_XXXXXXX_keymanager` command at the boot menu prompt multiple times until the nodes completely boot.

9. Boot the nodes:

`boot_ontap`

10. Wait for the replaced nodes to boot up.

If either node is in takeover mode, perform a giveback:

`storage failover giveback`

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

- a. View the broadcast domains:

`network port broadcast-domain show`

- b. Add any ports to a broadcast domain as needed.

[Adding or removing 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.

[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> |

### 2.7.8. Verifying 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.

### 2.7.9. Install the new licenses

Before the switchback operation, you must install licenses for the new controllers.

### Steps

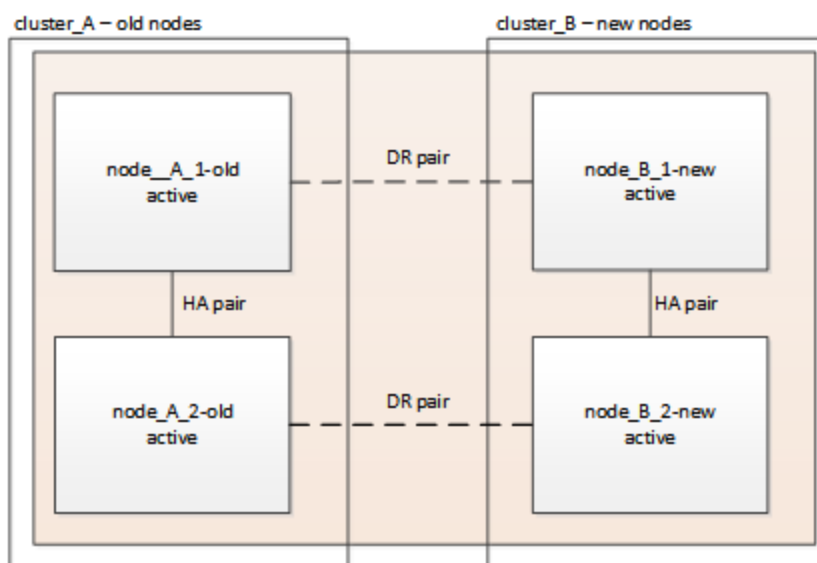
1. Install licenses for the new controller module.

## 2.8. Switching 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.

## 2.9. Checking 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:

### Listing 1. 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.

## 2.10. Upgrading 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 [Preparing for the upgrade](#).

As you perform the tasks, 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.

## 2.11. Sending 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.

## 2.12. 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 [Adding MetroCluster configurations](#) in *MetroCluster Tiebreaker Installation and Configuration*.

## Chapter 3. Upgrading 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.

### 3.1. 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 [Choosing an upgrade or refresh method](#) for additional procedures.

### 3.2. 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.
- 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), 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

### 3.3. Preparing 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. Start the automated controller replacement procedure from site A to replace the controllers at site B:

```
system controller replace start
```

The automated operation executes the prechecks. 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 are booted up. To bring the cluster back to a healthy state, you need to 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.

### 3.3.1. Gathering 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.
```

#### 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:

- `fcg adapter show -instance`
- `fcg 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
```

### 3.3.2. Removing the existing configuration from the Tiebreaker or other monitoring software

If the existing configuration is monitored with the MetroCluster Tiebreaker configuration 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 MetroCluster configurations](#) 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.

## 3.4. Replacing the old controllers and booting 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.

### 3.4.1. Preparing 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 [Preparing 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 [Lenovo Press](#)

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

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

In the following example, this is “e0d”.

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

3. Modify broadcast domains to remove VLAN and physical ports that need to be deleted:

```
broadcast-domain remove-ports -broadcast-domain broadcast-domain-name -ports node-name:port-id
```

Repeat this step for all VLAN and physical ports.

4. Remove any VLAN ports using cluster ports as member ports and interface groups using cluster ports as member ports.

- a. Delete VLAN ports:

```
network port vlan delete -node node-name -vlan-name portid-vlandid
```

For example:

```
network port vlan delete -node node1 -vlan-name e1c-80
```

- b. Remove physical ports from the interface groups:

```
network port ifgrp remove-port -node node-name -ifgrp interface-group-name -port portid
```

For example:

```
network port ifgrp remove-port -node node1 -ifgrp a1a -port e0d
```

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

```
network port broadcast-domain remove-ports -ip-space ip-space -broadcast-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..
```

- d. Modify interface group ports to use other physical ports as member as needed.:

```
ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id
```

5. Halt the nodes:

```
halt -inhibit-takeover true -node node-name
```

This step must be performed on both nodes.

### 3.4.2. Setting 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.

[ThinkSystem DG and DM Series](#)

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 [Lenovo Press](#) 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.

6. Power up the new nodes and press Ctrl-C when prompted to display the LOADER prompt.

### 3.4.3. Netbooting 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.

1. Access [Lenovo Data Center Support](#) to download the files used for performing the netboot of the system.
2. Download the appropriate ONTAP software 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.

| If the platform model is... | Then...                                                                                                                                                                                          |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| All systems                 | Your directory listing should contain a netboot folder with a kernel file: <code>ontap-version_image.tgz</code><br><br>You do not need to extract the <code>ontap-version_image.tgz</code> 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 by using this 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. Be sure to 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}
// end include reference
```

### 3.4.4. Clearing the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the configuration.

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. Restart the system by entering `bye`
5. Press Ctrl+C to enter the LOADER prompt.
6. At the LOADER prompt, launch the boot menu: `boot_ontap menu`
7. At the boot menu prompt, type the following command to clear the configuration: `wipeconfig`

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

8. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

### 3.4.5. Restoring the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site's usage.

#### Steps

1. In Maintenance mode configure the settings for any HBAs in the system:
  - a. Check the current settings of the ports: `ucadmin show`
  - b. Update the port settings as needed.

| If you have this type of HBA and desired mode... | Use this command...                                         |
|--------------------------------------------------|-------------------------------------------------------------|
| CNA FC                                           | <code>ucadmin modify -m fc -t initiator adapter-name</code> |
| CNA Ethernet                                     | <code>ucadmin modify -mode cna adapter-name</code>          |
| FC target                                        | <code>fcadmin config -t target adapter-name</code>          |
| FC initiator                                     | <code>fcadmin config -t initiator adapter-name</code>       |

2. Exit Maintenance mode:

`halt`

After you run the command, wait until the node stops at the LOADER prompt.

3. Boot the node back into Maintenance mode to enable the configuration changes to take effect:



```
boot_ontap maint
```

4. Verify the changes you made:

| If you have this type of HBA... | Use this command...       |
|---------------------------------|---------------------------|
| CNA                             | <code>ucadmin show</code> |
| FC                              | <code>fcadmin show</code> |

### 3.4.6. Reassigning 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-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 HOME     | OWNER                    | POOL  | SERIAL NUMBER | HOME                     | DR |
|---------------|--------------------------|-------|---------------|--------------------------|----|
| rr18:9.126L44 | node_B_1-old(4068741254) | Poo11 | PZHYN0MD      | node_B_1-old(4068741254) |    |
| rr18:9.126L49 | node_B_1-old(4068741254) | Poo11 | PPG3J5HA      | node_B_1-old(4068741254) |    |
| rr18:8.126L21 | node_B_1-old(4068741254) | Poo11 | PZHTDSZD      | node_B_1-old(4068741254) |    |
| rr18:8.126L2  | node_B_1-old(4068741254) | Poo10 | S0M1J2CF      | node_B_1-old(4068741254) |    |
| rr18:8.126L3  | node_B_1-old(4068741254) | Poo10 | S0M0CQM5      | node_B_1-old(4068741254) |    |
| rr18:9.126L27 | node_B_1-old(4068741254) | Poo10 | S0M1PSDW      | node_B_1-old(4068741254) |    |

- 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
```

- 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) Poo11 PZHYN0MD node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:9.126L49 node_B_1-new(1574774970) Poo11 PPG3J5HA node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:8.126L21 node_B_1-new(1574774970) Poo11 PZHTDSZD node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:8.126L2 node_B_1-new(1574774970) Poo10 S0M1J2CF node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:9.126L29 node_B_1-new(1574774970) Poo10 S0M0CQM5 node_B_1-new(1574774970)
node_B_1-new(1574774970)
rr18:8.126L1 node_B_1-new(1574774970) Poo10 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).

### 3.4.7. Booting up the new controllers

You must reboot the controllers from the boot menu to update the controller flash image. Additional steps are required if encryption is configured.

You can reconfigure VLANs and interface groups. If required, manually modify the ports for the cluster LIFs and broadcast domain details before resuming the operation by using the **system controller replace resume** command.

#### About this task

This task must be performed on all the new controllers.

#### Steps

1. Halt the node:

**halt**

2. If external key manager is configured, set the related bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address
setenv bootarg.kmip.init.netmask netmask
setenv bootarg.kmip.init.gateway gateway-address
setenv bootarg.kmip.init.interface interface-id
```

3. Display the boot menu:

```
boot_ontap menu
```

4. If root encryption is used, select the boot menu option for your key management configuration.

| If you are using...     | Select this boot menu option...                                                                                            |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option "10"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option "11"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |

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"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option "11"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |

Depending on the key manager setting, perform the recovery procedure by selecting option "10" or option "11", followed by option "6" at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option "1" (normal boot).

9. Boot the nodes:

`boot_ontap`

10. Wait for the replaced nodes to boot up.

If either node is in takeover mode, perform a giveback using the `storage failover giveback` command.

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

- a. View the broadcast domains:

`network port broadcast-domain show`

- b. Add any ports to a broadcast domain as needed.

[Adding or removing ports from a broadcast domain](#)

- c. Add the physical port that will host the intercluster LIFs to the corresponding broadcast domain.
- d. Modify intercluster LIFs to use the new physical port as home port.
- e. After the intercluster LIFs are up, check the cluster peer status and re-establish cluster peering as needed.

You may need to reconfigure cluster peering.

### Creating a cluster peer relationship

- f. Recreate VLANs and interface groups as needed.

VLAN and interface group membership might be different than that of the old node.

### Creating a VLAN

### Combining physical ports to create interface groups

- g. Verify that the partner cluster is reachable and that the configuration is successfully resynchronized on the partner cluster:

```
metrocluster switchback -simulate true
```

12. If encryption is used, restore the keys using the correct command for your key management configuration.

| If you are using...     | Use this command...                                                                                                                                                                                                                                                      |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | <pre>security key-manager onboard sync</pre> <p>For more information, see <a href="#">Restoring onboard key management encryption keys</a>.</p>                                                                                                                          |
| External key management | <pre>security key-manager external restore<br/>-vserver SVM -node node -key-server<br/>host_name IP_address:port -key-id key_id<br/>-key-tag key_tag node-name</pre> <p>For more information, see <a href="#">Restoring 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
```

## 3.5. Completing 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.

### 3.5.1. 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 [Adding MetroCluster configurations](#).

## Chapter 4. Upgrading controllers in a MetroCluster IP configuration 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.

### 4.1. 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.
- Platform models with internal shelves are not supported.
- 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.

### 4.2. About this task

- 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 IP switches must be running a supported firmware version.
- 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 [Lenovo Press](#).

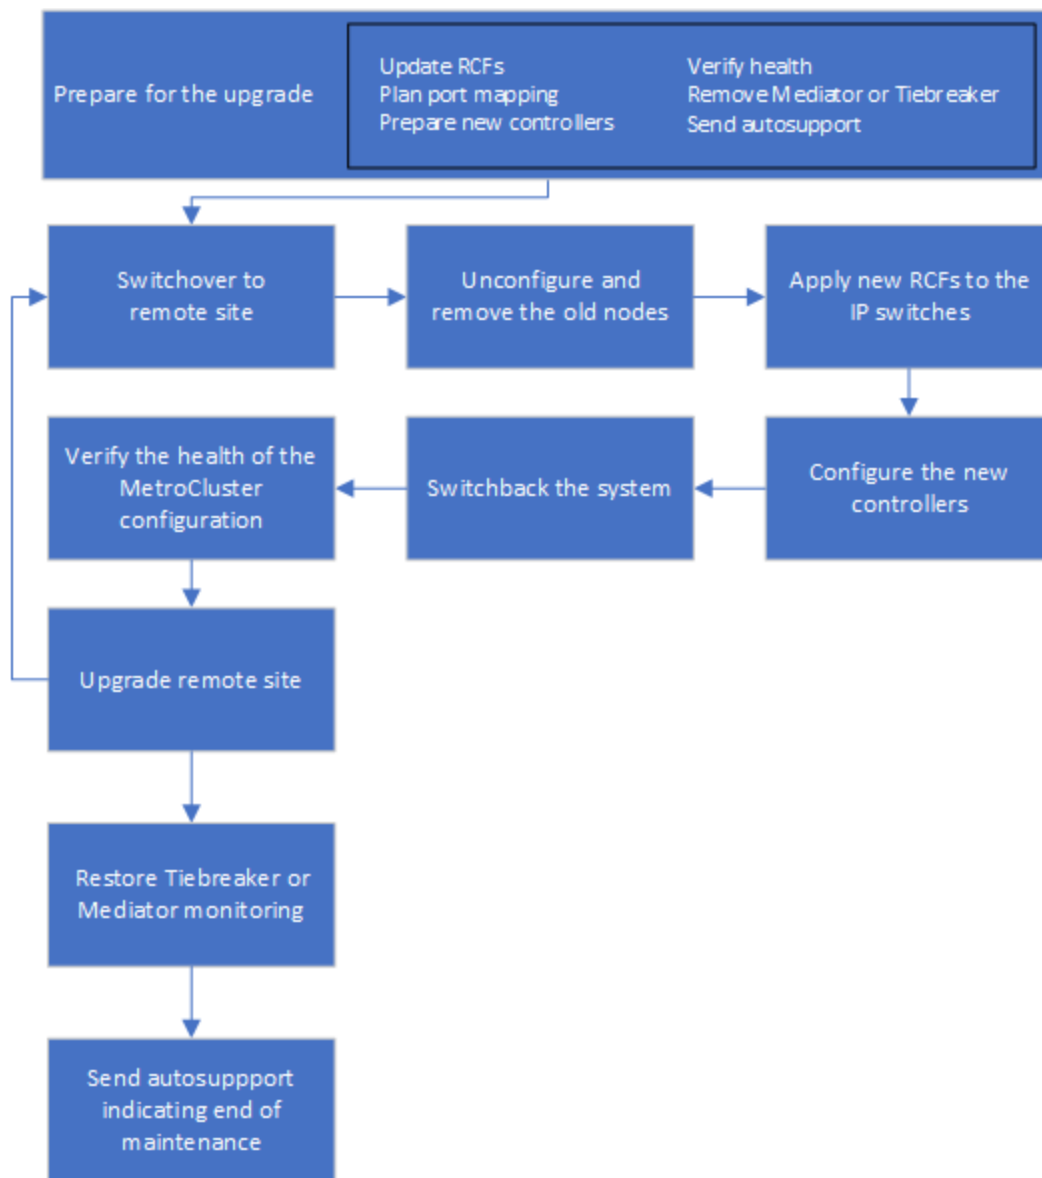
- You will reuse the IP addresses, netmasks, and gateways of the original platforms on the new platforms.
- The following example names are used in this procedure:
  - site\_A
    - Before upgrade:



- 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

### **4.3. Workflow for upgrading controllers in an MetroCluster IP configuration**

You can use the workflow diagram to help you plan the upgrade tasks.



## 4.4. Preparing 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.

### 4.4.1. Updating the MetroCluster switch RCF files before upgrading controllers

Depending on the old platform models, or if switch configuration is not on the minimum version, or if you want to change VLAN IDs used by the back-end MetroCluster connections, you must update the switch RCF files before you begin the platform upgrade procedure.

#### About this task

You must update the RCF file in the following scenarios:

- For certain platform models, the switches must be using a supported VLAN ID for the back-end

MetroCluster IP connections. If the old or new platform models are in the following table, **and not** using a supported VLAN ID, you must update the switch RCF files.



The local cluster connections can use any VLAN, they do not need to be in the given range.

| Platform model (old or new)                               | Supported VLAN IDs                                                                                                      |
|-----------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>DM7100F</li> </ul> | <ul style="list-style-type: none"> <li>10</li> <li>20</li> <li>Any value in the range 101 to 4096 inclusive.</li> </ul> |

- The switch configuration was not configured with minimum supported RCF version:

| Switch model       | Required RCF file version |
|--------------------|---------------------------|
| Cisco 3132Q-V      | 1.7 or later              |
| Cisco 3232C        | 1.7 or later              |
| Broadcom BES-53248 | 1.3 or later              |

- You want to change the VLAN configuration.

The VLAN ID range is 101 to 4096 inclusive.

The switches at site\_A will be upgraded when the controllers on site\_A are upgraded.

## Steps

- Prepare the IP switches for the application of the new RCF files.

Follow the steps in the section for your switch vendor from the [MetroCluster IP installation and configuration](#).

- [Resetting the Broadcom IP switch to factory defaults](#)
- [Resetting the Cisco IP switch to factory defaults](#)

- Download and install the RCF files.

Follow the steps in the [MetroCluster IP installation and configuration](#).

- [Downloading and installing the Broadcom RCF files](#)
- [Downloading and installing the Cisco IP RCF files](#)

#### 4.4.2. Mapping 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, which will allow 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 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 upgrade, you must recreate 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 [Lenovo Press](#).

2. Plan your port usage and fill in the following tables for reference for each of the new nodes.

You will refer to the table as you carry out the upgrade procedure.

| LIF       | node_A_1-old |          |                   | node_A_1-new |          |                   |
|-----------|--------------|----------|-------------------|--------------|----------|-------------------|
|           | Ports        | IPspaces | Broadcast domains | Ports        | IPspaces | Broadcast domains |
| Cluster 1 |              |          |                   |              |          |                   |
| Cluster 2 |              |          |                   |              |          |                   |
| Cluster 3 |              |          |                   |              |          |                   |

|                    |  |  |  |  |  |  |
|--------------------|--|--|--|--|--|--|
| Cluster 4          |  |  |  |  |  |  |
| Node management    |  |  |  |  |  |  |
| Cluster management |  |  |  |  |  |  |
| Data 1             |  |  |  |  |  |  |
| Data 2             |  |  |  |  |  |  |
| Data 3             |  |  |  |  |  |  |
| Data 4             |  |  |  |  |  |  |
| SAN                |  |  |  |  |  |  |
| Intercluster port  |  |  |  |  |  |  |

### 4.4.3. Netbooting 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 [Lenovo Data Center Support](#) 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 Lenovo Data Center Support 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.
  - d. At the LOADER prompt, configure the netboot connection for a management LIF:

|                               |                |
|-------------------------------|----------------|
| <b>If IP addressing is...</b> | <b>Then...</b> |
|-------------------------------|----------------|

|        |                                                                                                               |
|--------|---------------------------------------------------------------------------------------------------------------|
| DHCP   | Configure the automatic connection:<br><br><code>ifconfig e0M -auto</code>                                    |
| Static | Configure the manual connection:<br><br><code>ifconfig e0M -addr=ip_addr -mask=netmask<br/>-gw=gateway</code> |

e. Perform the netboot.

| If the platform model is... | Then...                                                                                                           |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------|
| All systems                 | <code>netboot<br/>http://_web_server_ip/path_to_web-<br/>accessible_directory/ontap-<br/>version_image.tgz</code> |

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. Be sure to 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}
```

#### 4.4.4. Clearing the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the configuration.

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. Restart the system by entering `bye`
5. Press Ctrl+C to enter the LOADER prompt.
6. At the LOADER prompt, launch the boot menu: `boot_ontap menu`
7. At the boot menu prompt, type the following command to clear the configuration: `wipeconfig`

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

8. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

#### 4.4.5. Verifying MetroCluster health before site upgrade

You must verify the health and connectivity of the MetroCluster configuration prior to 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 is 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. 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
```

#### 4.4.6. Gathering 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.

You should 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 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::~*
```

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  | 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 |

- If the MetroCluster nodes are in a SAN configuration, collect the relevant information.

You should gather the output of the following commands:

- `fcg adapter show -instance`
- `fcg interface show -instance`
- `iscsi interface show`
- `ucadmin show`

- 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 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.

```

**4.4.7. Removing 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> in the <i>MetroCluster Tiebreaker Installation and Configuration Guide</i> |
| Mediator            | Issue the following command from the ONTAP prompt:<br><br><code>metrocluster configuration-settings mediator remove</code>      |

**4.4.8. Sending a custom AutoSupport message prior to maintenance**

Before performing the maintenance, you should issue an AutoSupport message to notify Lenovo 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.

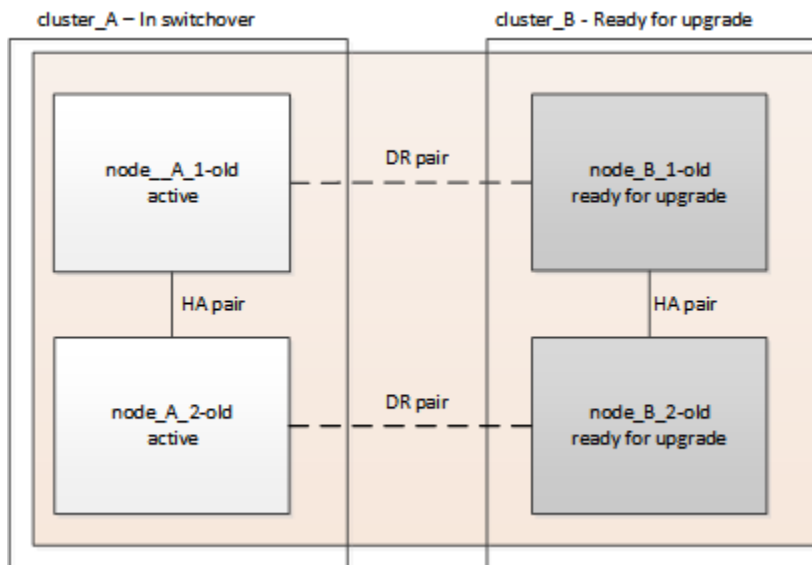
## 4.5. Switching over the MetroCluster configuration

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.



### 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 controller upgrade.

## 4.6. Removing interface configurations and uninstalling the old controllers

You must move data LIFs to a common port, remove VLANs and interface groups on the old controllers and then physically uninstall the controllers.

### About this task

- These steps are performed on the old controllers (node\_B\_1-old, node\_B\_2-old).
- See the information you gathered in [Mapping ports from the old nodes to the new nodes](#).

### Steps

1. Boot the old nodes and 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 [Lenovo Press](#)

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

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

In the following example, this is "e0d".

For example:

```
network interface modify -vserver vs0 -lif datalif1 -home-port e0d
```

3. 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 -ip-space ip-space -broadcast-domain broadcast-domain-name -ports nodename:portname,nodename:portname,..
```

- d. Modify interface group ports to use other physical ports as member as needed.:

```
ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id
```

4. Halt the nodes to the LOADER prompt:

```
halt -inhibit-takeover true
```

5. 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.

6. Gather the bootarg values:

```
printenv
```

7. 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.
8. Disconnect the power cables from node\_B\_1-old and node\_B\_2-old.
9. Remove the node\_B\_1-old and node\_B\_2-old controllers from the rack.

## 4.7. Updating 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 RCF files.

Follow the steps in the procedure for your switch vendor:

#### [MetroCluster IP installation and configuration](#)

- [Resetting the Broadcom IP switch to factory defaults](#)
- [Resetting the Cisco IP switch to factory defaults](#)

2. Download and install the RCF files.

Follow the steps in the section for your switch vendor from the [MetroCluster IP installation and configuration](#).

- [Downloading and installing the Broadcom RCF files](#)
- [Downloading and installing the Cisco IP RCF files](#)

## 4.8. Configuring 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.

### 4.8.1. Setting 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.

#### ThinkSystem DG and DM Series

4. Cable the controllers to the IP switches as described in [MetroCluster IP installation and configuration](#).
1. Power up the new nodes and boot them to Maintenance mode.

### 4.8.2. Restoring the HBA configuration

Depending on the presence and configuration of HBA cards in the controller module, you need to configure them correctly for your site's usage.

#### Steps

1. In Maintenance mode configure the settings for any HBAs in the system:
  - a. Check the current settings of the ports:

```
ucadmin show
```

- b. Update the port settings as needed.

| If you have this type of HBA and desired mode... | Use this command...                                         |
|--------------------------------------------------|-------------------------------------------------------------|
| CNA FC                                           | <code>ucadmin modify -m fc -t initiator adapter-name</code> |
| CNA Ethernet                                     | <code>ucadmin modify -mode cna adapter-name</code>          |
| FC target                                        | <code>fcadmin config -t target adapter-name</code>          |
| FC initiator                                     | <code>fcadmin config -t initiator adapter-name</code>       |

2. Exit Maintenance mode:

```
halt
```

After you run the command, wait until the node stops at the LOADER prompt.

3. Boot the node back into Maintenance mode to enable the configuration changes to take effect:

```
boot_ontap maint
```



4. Verify the changes you made:

| If you have this type of HBA... | Use this command...       |
|---------------------------------|---------------------------|
| CNA                             | <code>ucadmin show</code> |
| FC                              | <code>fcadmin show</code> |

### 4.8.3. Setting 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
```

### 4.8.4. Setting 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 [Gathering information before the upgrade](#).

#### Steps

1. If the nodes being upgraded are DM7100H, or DM7100F models, set the following bootargs at the LOADER prompt:

```
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 the default VLANs, the vlan-id is not necessary.

The following commands set 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 commands set 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 shows the commands for node\_B\_1-new when the default VLAN is used:

```
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 shows the commands for node\_B\_2-new when the default VLAN is used:

```
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. If the nodes being upgraded are not systems listed in the previous step, at the LOADER prompt for each of the surviving nodes, set the following bootargs with local\_IP/mask:

```
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
```

```
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
```

The following commands set the values for node\_B\_1-new:

```
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 commands set the values for node\_B\_2-new:

```
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
```

3. 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
```

4. If the original systems were configured for ADP, at each of the replacement nodes' LOADER prompt, enable ADP:

```
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
```

#### 4.8.5. Reassigning 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.



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
```

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-old(537403322) Pool1 PZHYN0MD node_B_1-old(537403322)
node_B_1-old(537403322)
prod4-rk18:9.126L49 node_B_1-old(537403322) Pool1 PPG3J5HA node_B_1-old(537403322)
node_B_1-old(537403322)
prod4-rk18:8.126L21 node_B_1-old(537403322) Pool1 PZHTDSZD node_B_1-old(537403322)
node_B_1-old(537403322)
prod2-rk18:8.126L2 node_B_1-old(537403322) Pool10 S0M1J2CF node_B_1-old(537403322)
node_B_1-old(537403322)
prod2-rk18:8.126L3 node_B_1-old(537403322) Pool10 S0M0CQM5 node_B_1-old(537403322)
node_B_1-old(537403322)
prod1-rk18:9.126L27 node_B_1-old(537403322) Pool10 S0M1PSDW node_B_1-old(537403322)
node_B_1-old(537403322)
.
.
.
```

3. Reassign the root aggregate disks on the drive shelves to the new controllers.

| If you are using ADP... | Then use this command...                                                 |
|-------------------------|--------------------------------------------------------------------------|
| Yes                     | <code>disk reassign -s old-sysid -d new-sysid -r dr-partner-sysid</code> |
| No                      | <code>disk reassign -s old-sysid -d new-sysid</code>                     |

4. Reassign the root aggregate disks on the drive shelves to the new controllers:

```
disk reassign -s old-sysid -d new-sysid
```

The following example shows reassignment of drives in a non-ADP configuration:

```
*> disk reassign -s 537403322 -d 1574774970
Partner node must not be in Takeover mode during disk reassignment from maintenance mode.
Serious problems could result!!
Do not proceed with reassignment if the partner is in takeover mode. Abort reassignment (y/n)? n

After the node becomes operational, you must perform a takeover and giveback of the HA partner
node to ensure disk reassignment is successful.
Do you want to continue (y/n)? y
Disk ownership will be updated on all disks previously belonging to Filer with sysid 537403322.
Do you want to continue (y/n)? y
```

5. Verify that the disks of the root aggregate are properly 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-new(537097247) Pool11 PZHYN0MD node_B_1-new(537097247)
node_B_1-new(537097247)
prod04-rk18:9.126L49 node_B_1-new(537097247) Pool11 PPG3J5HA node_B_1-new(537097247)
node_B_1-new(537097247)
prod04-rk18:8.126L21 node_B_1-new(537097247) Pool11 PZHTDSZD node_B_1-new(537097247)
node_B_1-new(537097247)
prod02-rk18:8.126L2 node_B_1-new(537097247) Pool10 S0M1J2CF node_B_1-new(537097247)
node_B_1-new(537097247)
prod02-rk18:9.126L29 node_B_1-new(537097247) Pool10 S0M0CQM5 node_B_1-new(537097247)
node_B_1-new(537097247)
prod01-rk18:8.126L1 node_B_1-new(537097247) Pool10 S0M1PSDW node_B_1-new(537097247)
node_B_1-new(537097247)
::>
::> aggr status
 Aggr State Status Options
aggr0_node_B_1 online raid_dp, aggr root, nosnap=on,
 mirrored
 fast zeroed
 64-bit
```

### 4.8.6. Booting 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 <b>10</b><br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option <b>11</b><br><br>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 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. On LOADER, double-check the bootarg values and update the values as needed.

Use the steps in [Setting the MetroCluster IP bootarg variables](#).

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 again for your key management configuration.

| If you are using...     | Select this boot menu option...                                                                                                 |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------|
| Onboard key management  | Option <b>10</b><br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration. |
| External key management | Option "11"<br><br>Follow the prompts to provide the required inputs to recover and restore the key-manager configuration.      |

Depending on the key manager setting, perform the recovery procedure by selecting option "10" or option "11", followed by option **6** at the first boot menu prompt. To boot the nodes completely, you might need to repeat the recovery procedure continued by option "1" (normal boot).

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. 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><br><br>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><br><br>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](#)

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

#### 4.8.7. Verifying and restoring 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 [Mapping 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
```

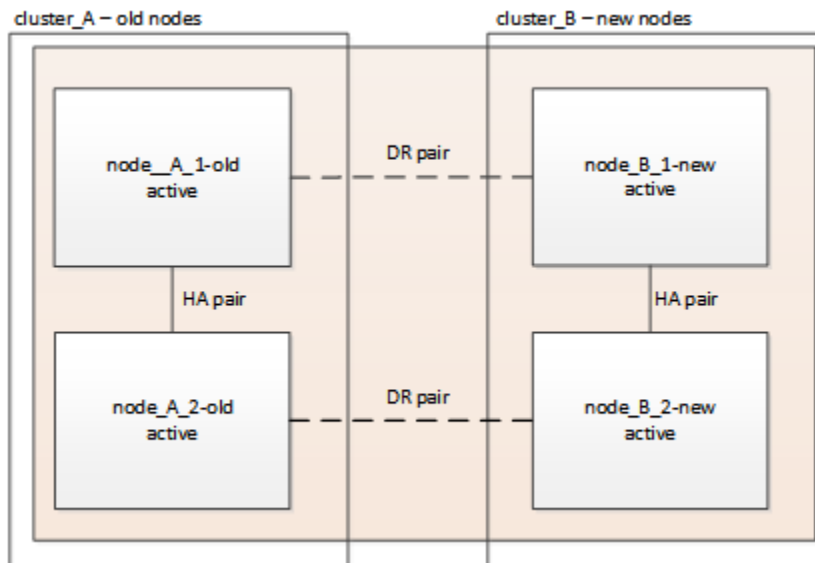
2. Revert the interfaces to their home node:

```
network interface revert * -vserver vserver-name
```

Perform this step on all SVMs as required.

## 4.9. Switching 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 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.

## 4.10. Checking 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
```

## 4.11. Upgrading the nodes on cluster\_A

You must repeat the upgrade tasks on cluster\_A.

### Steps

1. Repeat the steps to upgrade the nodes on cluster\_A, beginning with [Preparing for the upgrade](#).

As you perform the tasks, 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.

## 4.12. Restoring 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> .                                          |
| Mediator                 | <a href="#">Configuring the ONTAP Mediator service from a MetroCluster IP configuration</a> . |
| Third-party applications | Refer to the product documentation.                                                           |

## 4.13. Sending 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.

## Chapter 5. 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

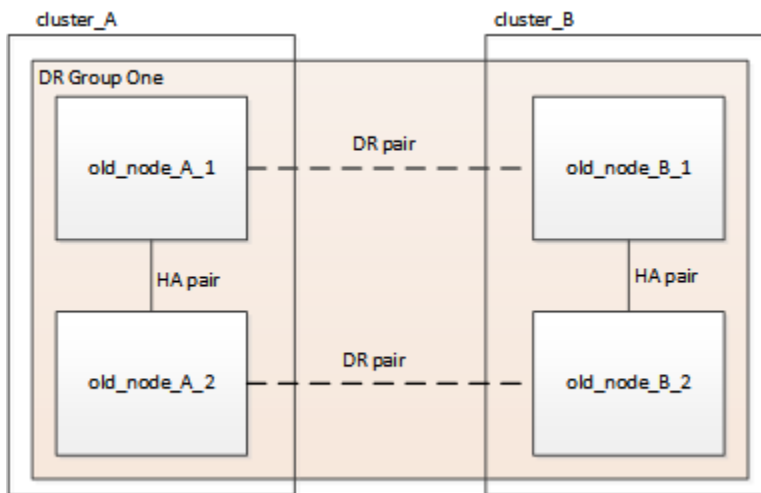
References to "old nodes" mean the nodes that you intend to replace.

- 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](#).

### Steps

1. Gather information from the old nodes.

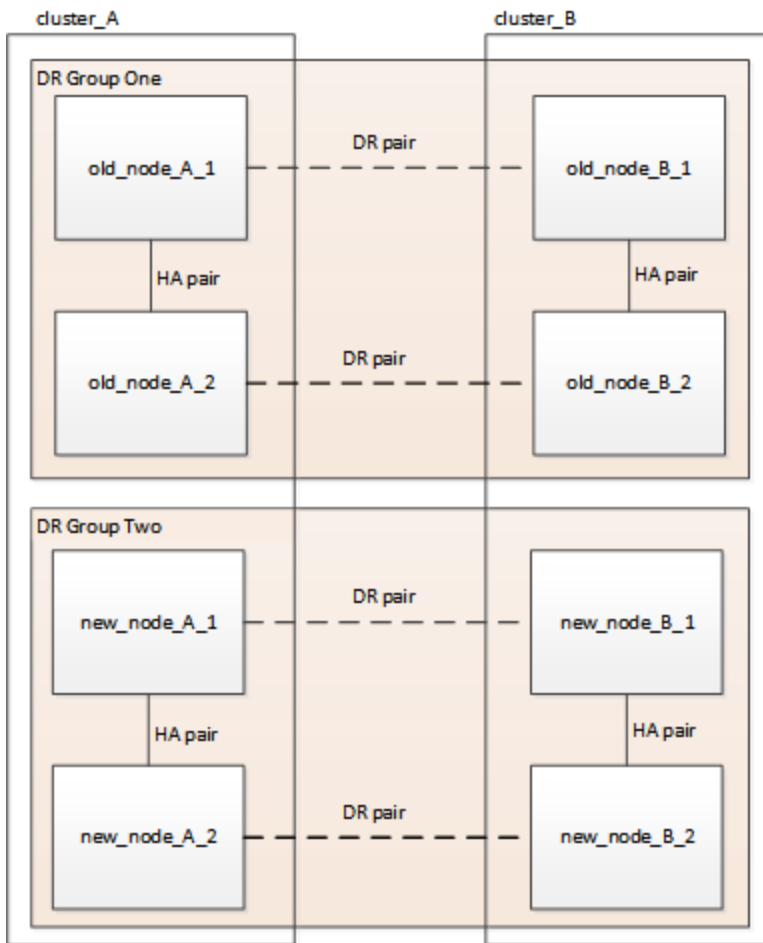
At this stage, the four-node configuration appears as shown in the following image:



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 [Moving a metadata volume in MetroCluster configurations](#).

4. Move the data from the old nodes to new nodes using the following procedures.
  - a. Create an aggregate and move volumes to the new nodes.

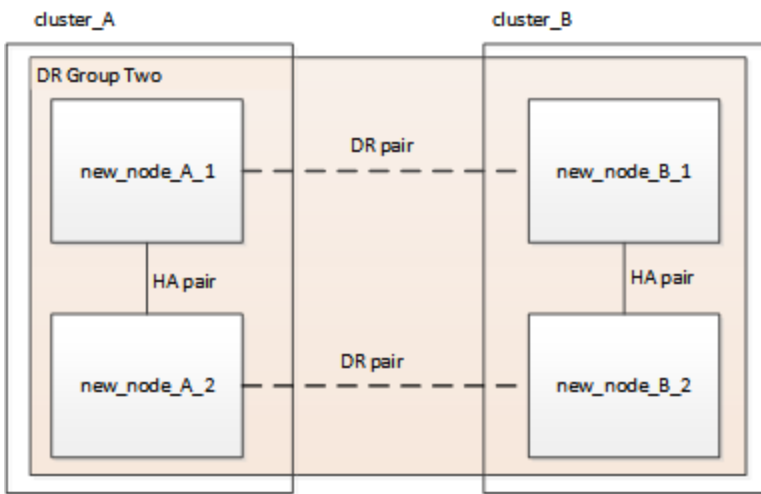


You might choose to mirror the aggregate when or after it is created.

- b. Move non-SAN data LIFs and cluster management LIFs to the new nodes.
  - c. Delete SAN LIFs 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:





## Chapter 6. Refreshing 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.

### 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.

If you refresh both DR groups, you must refresh one DR group 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.



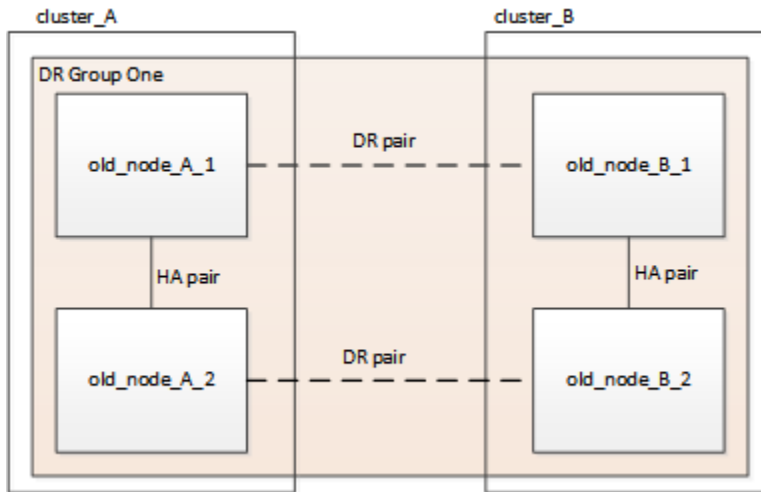
If you refresh both DR groups, the platform combination might not be supported after you refresh the first DR group. You must refresh both DR groups to achieve a supported eight-node configuration.

- 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.

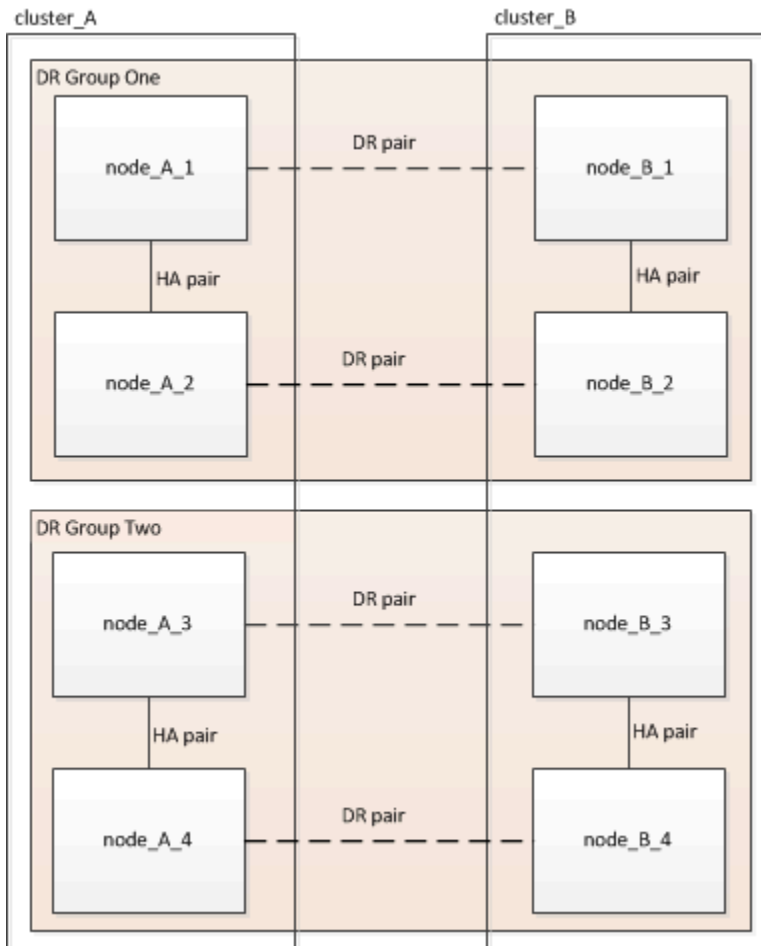
### Steps

1. 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:



2. 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. 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.

3. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

| If you are using...      | Use this procedure...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tiebreaker               | <p>1. Use the Tiebreaker CLI <code>monitor remove</code> command to remove the MetroCluster configuration.</p> <p>In the following example, “cluster_A” is removed from the software:</p> <pre>NetApp MetroCluster Tiebreaker :&gt; monitor remove -monitor-name cluster_A Successfully removed monitor from NetApp MetroCluster Tiebreaker software.</pre> <p>2. Confirm that the MetroCluster configuration is removed correctly by using the Tiebreaker CLI <code>monitor show -status</code> command.</p> <pre>NetApp MetroCluster Tiebreaker :&gt; monitor show -status</pre> |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

4. 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:

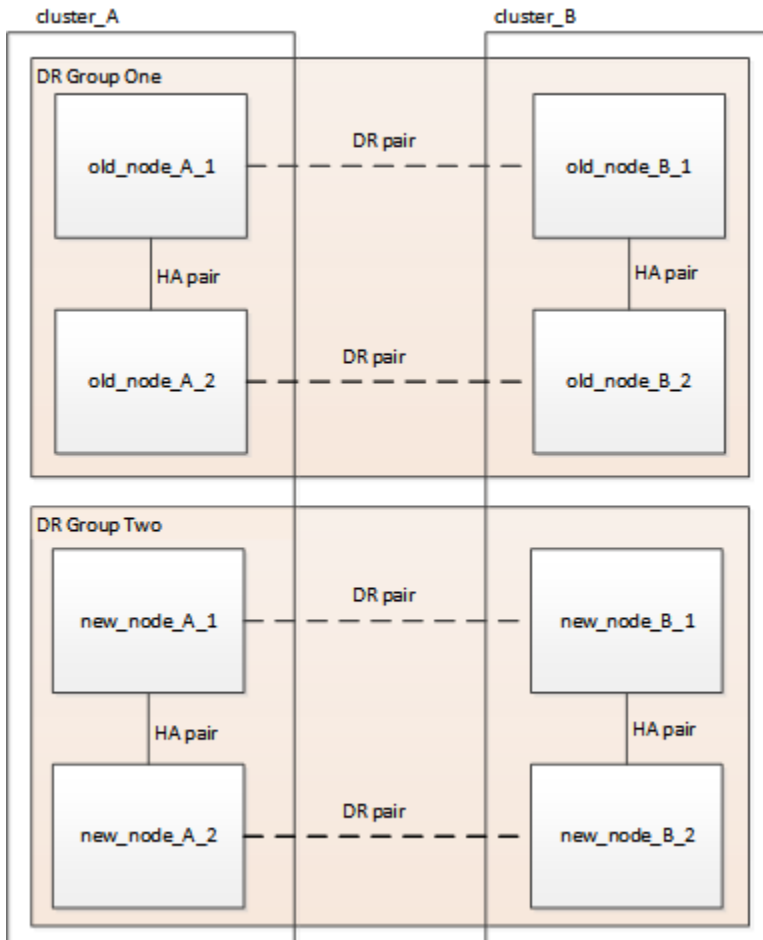
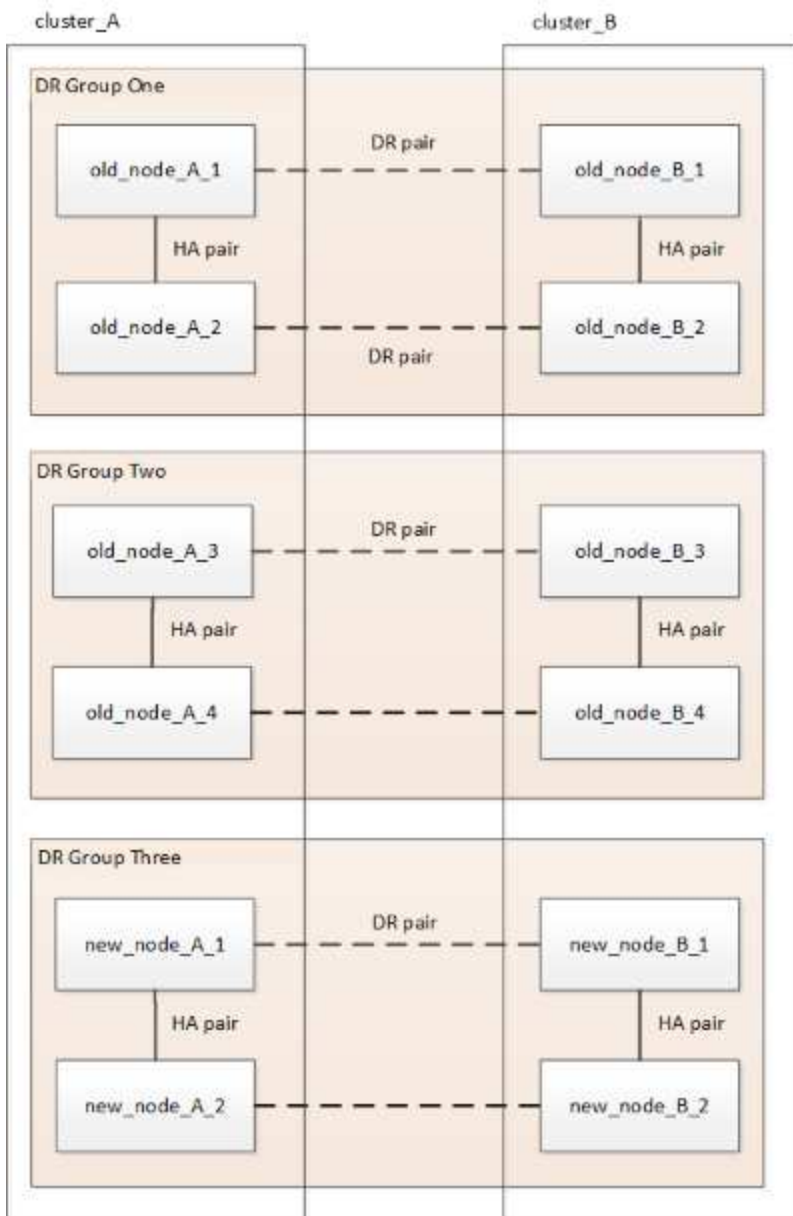


Figure 1. Temporary eight-node configuration



**Figure 2. Temporary twelve-node configuration**

5. Confirm that takeover is possible and the nodes are connected by running the following command on both clusters:

`storage failover show`

```
cluster_A::> storage failover show
```

| Node      | Partner   | Takeover Possible | State Description      |
|-----------|-----------|-------------------|------------------------|
| Node_FC_1 | Node_FC_2 | true              | Connected to Node_FC_2 |
| Node_FC_2 | Node_FC_1 | true              | Connected to Node_FC_1 |
| Node_IP_1 | Node_IP_2 | true              | Connected to Node_IP_2 |
| Node_IP_2 | Node_IP_1 | true              | Connected to Node_IP_1 |

6. Move the CRS volumes.

Perform the steps in [Moving a metadata volume in MetroCluster configurations](#).

7. Move the data from the old nodes to the new nodes. by using the following procedures in [Refreshing a four-node MetroCluster IP configuration](#)

- a. Create an aggregate and moving volumes to the new nodes.



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

- b. Move non-SAN data LIFs and cluster management LIFs to the new nodes.

8. Modify the IP address for the cluster peer of the transitioned nodes for each cluster:

- a. Identify the cluster\_A peer by using the `cluster peer show` command:

```
cluster_A::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication

cluster_B 1-80-000011 Unavailable absent
```

- i. Modify the cluster\_A peer IP address:

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

- b. Identify the cluster\_B peer by using the `cluster peer show` command:

```
cluster_B::> cluster peer show
Peer Cluster Name Cluster Serial Number Availability Authentication

cluster_A 1-80-000011 Unavailable absent
```

- i. Modify the cluster\_B peer IP address:

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

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

- i. Verify that the IP address is updated for each cluster by using the `cluster peer show -instance` command.

The `Remote Intercluster Addresses` field in the following examples displays the updated IP address.

Example for cluster\_A:

```

cluster_A::> cluster peer show -instance

Peer Cluster Name: cluster_B
 Remote Intercluster Addresses: 172.21.178.204, 172.21.178.212
 Availability of the Remote Cluster: Available
 Remote Cluster Name: cluster_B
 Active IP Addresses: 172.21.178.212, 172.21.178.204
 Cluster Serial Number: 1-80-000011
 Remote Cluster Nodes: node_B_3-IP,
 node_B_4-IP
 Remote Cluster Health: true
 Unreachable Local Nodes: -
 Address Family of Relationship: ipv4
 Authentication Status Administrative: use-authentication
 Authentication Status Operational: ok
 Last Update Time: 4/20/2023 18:23:53
 IPspace for the Relationship: Default
 Proposed Setting for Encryption of Inter-Cluster Communication: -
 Encryption Protocol For Inter-Cluster Communication: tls-psk
 Algorithm By Which the PSK Was Derived: jpake

cluster_A::>

```

#### Example for cluster\_B

```

cluster_B::> cluster peer show -instance


Peer Cluster Name: cluster_A
 Remote Intercluster Addresses: 172.21.178.188, 172.21.178.196 <<<<<<<< Should
reflect the modified address
 Availability of the Remote Cluster: Available
 Remote Cluster Name: cluster_A
 Active IP Addresses: 172.21.178.196, 172.21.178.188
 Cluster Serial Number: 1-80-000011
 Remote Cluster Nodes: node_A_3-IP,
 node_A_4-IP
 Remote Cluster Health: true
 Unreachable Local Nodes: -
 Address Family of Relationship: ipv4
 Authentication Status Administrative: use-authentication
 Authentication Status Operational: ok
 Last Update Time: 4/20/2023 18:23:53
 IPspace for the Relationship: Default
 Proposed Setting for Encryption of Inter-Cluster Communication: -
 Encryption Protocol For Inter-Cluster Communication: tls-psk
 Algorithm By Which the PSK Was Derived: jpake

cluster_B::>

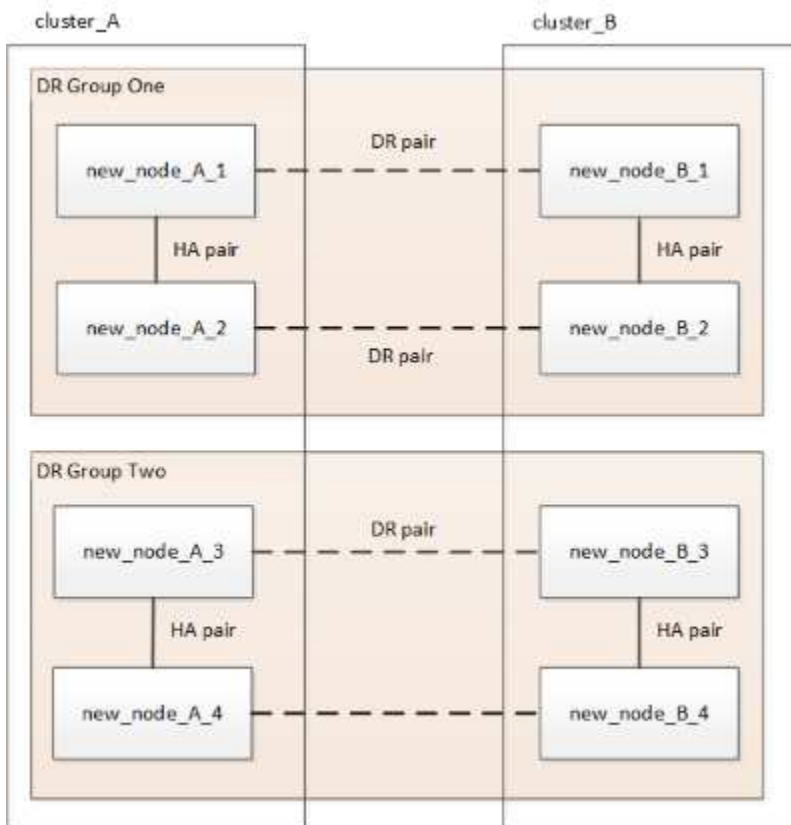
```

9. Follow the steps in [Removing a Disaster Recovery group](#) to remove the old DR group.
10. If you want to refresh both DR groups in an eight-node configuration, you must 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**

11. 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
```

12. Restore monitoring if necessary, using the procedure for your configuration.

| If you are using... | Use this procedure                                                                                                                                         |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Tiebreaker          | <a href="#">Adding MetroCluster configurations</a> in the <i>MetroCluster Tiebreaker Installation and Configuration</i> .                                  |
| Mediator            | <a href="#">Configuring the ONTAP Mediator service from a MetroCluster IP configuration</a> in the <i>MetroCluster IP Installation and Configuration</i> . |

13. 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.

## Chapter 7. Expand a four-node MetroCluster FC configuration to an eight-node configuration

### 7.1. 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.

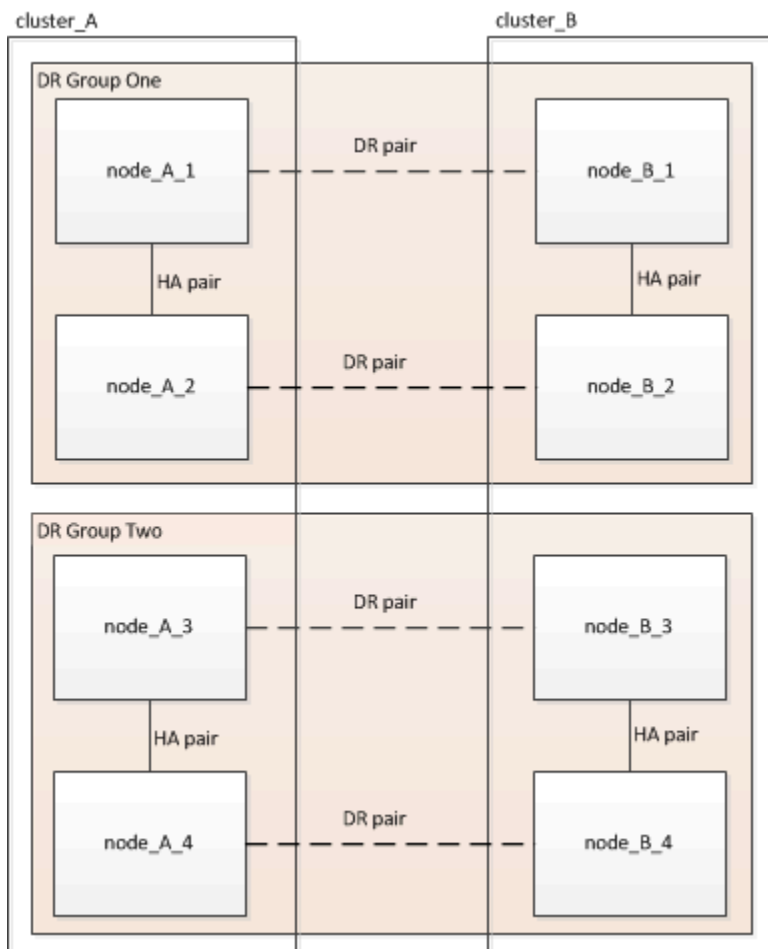
This procedure is not supported on earlier versions of ONTAP or in MetroCluster IP configurations.

- 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.

### 7.1.1. Supported platform combinations when adding a second DR group

The following table shows the supported platform combinations for eight-node MetroCluster FC configurations.



- All nodes in the MetroCluster configuration must be running the same ONTAP version. For example, if you have an eight-node configuration, all eight nodes must be running the same ONTAP version.
- The combinations in this table apply only to regular or permanent eight-node configurations.
- The platform combinations in this table **do not** apply if you are using the the transition or refresh procedures.
- All nodes in one DR group must be of the same type DR and configuration.

|                         |         | 8-node DR group 2 |         |         |         |
|-------------------------|---------|-------------------|---------|---------|---------|
|                         |         | DM7000H           | DM7000F | DM7100H | DM7100F |
| 8-node<br>DR group<br>1 | DM7000H |                   |         |         |         |
|                         | DM7000F |                   |         |         |         |
|                         | DM7100H |                   |         |         |         |
|                         | DM7100F |                   |         |         |         |

## 7.2. 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.

## 7.3. 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.

## 7.4. Verifying the health of the MetroCluster configuration

You should check the health of the MetroCluster configuration to verify proper operation.

### Steps

1. Check that the MetroCluster is configured and in normal mode on each cluster:

```
metrocluster show
```

```
cluster_A::> metrocluster show
Cluster Entry Name State

Local: cluster_A Configuration state configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-disaster
Remote: cluster_B Configuration state configured
 Mode normal
 AUSO Failure Domain auso-on-cluster-disaster
```

2. Check that mirroring is enabled on each node:

```
metrocluster node show
```

```
cluster_A::> metrocluster node show
DR
Group Cluster Node Configuration State DR
Mirroring Mode

1 cluster_A
 node_A_1 configured enabled normal
 cluster_B
 node_B_1 configured enabled normal
2 entries were displayed.
```

3. Check that the MetroCluster components are healthy:

`metrocluster check run`

```
cluster_A::> metrocluster check run
```

```
Last Checked On: 10/1/2014 16:03:37
```

| Component          | Result |
|--------------------|--------|
| nodes              | ok     |
| lifs               | ok     |
| config-replication | ok     |
| aggregates         | ok     |

4 entries were displayed.

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results.

To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

4. Check that there are no health alerts:

`system health alert show`

5. Simulate a switchover operation:

- a. From any node's prompt, change to the advanced privilege level:

`set -privilege advanced`

You need to respond with **y** when prompted to continue into advanced mode and see the advanced mode prompt (\*>).

- b. Perform the switchover operation with the -simulate parameter:

`metrocluster switchover -simulate`

- c. Return to the admin privilege level:

`set -privilege admin`

## 7.5. Sending a custom AutoSupport message prior to adding nodes to the MetroCluster configuration

You should issue an AutoSupport message to notify Lenovo 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.

## 7.6. Recable and zone a switch fabric for the new nodes

### 7.6.1. 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.

### 7.6.2. 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. 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.

## 7.7. Configure ONTAP on the new controllers

### 7.7.1. Clearing the configuration on a controller module

Before using a new controller module in the MetroCluster configuration, you must clear the configuration.

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. Restart the system by entering `bye`
5. Press Ctrl+C to enter the LOADER prompt.
6. At the LOADER prompt, launch the boot menu: `boot_ontap menu`
7. At the boot menu prompt, type the following command to clear the configuration: `wipeconfig`

Respond `yes` to the confirmation prompt.

The node reboots and the boot menu is displayed again.

8. At the boot menu, select option **5** to boot the system into Maintenance mode.

Respond `yes` to the confirmation prompt.

### 7.7.2. Assigning disk ownership in AFA systems

If you are using AFA 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.

### 7.7.3. Assigning disk ownership in non-AFA systems

If the MetroCluster nodes do not have the disks correctly assigned, or if you are using DM600S 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 DM600s 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 DM600s disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

| Assign these disks in the drawer... | To this node and pool...              |
|-------------------------------------|---------------------------------------|
| 0 - 2                               | Local node's pool 0                   |
| 3 - 5                               | HA partner node's pool 0              |
| 6 - 8                               | DR partner of the local node's pool 1 |
| 9 - 11                              | DR partner of the HA partner's pool 1 |

This disk assignment pattern ensures that an aggregate is minimally affected in case a drawer goes offline.

### Steps

1. If you have not done so, boot each system into Maintenance mode.
2. Assign the disk shelves to the nodes located at the first site (site A):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node, systematically assign the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

If storage controller Controller\_A\_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf1 -p 0
*> disk assign -shelf FC_switch_A_1:1-4.shelf2 -p 0

*> disk assign -shelf FC_switch_B_1:1-4.shelf1 -p 1
*> disk assign -shelf FC_switch_B_1:1-4.shelf2 -p 1
```

- b. Repeat the process for the second node at the local site, systematically assigning the local disk shelves to pool 0 and the remote disk shelves to pool 1:

```
disk assign -shelf local-switch-name:shelf-name.port -p pool
```

If storage controller Controller\_A\_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1

*> disk assign -shelf FC_switch_A_1:1-4.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-4.shelf4 -p 1
```

3. Assign the disk shelves to the nodes located at the second site (site B):

Disk shelves at the same site as the node are assigned to pool 0 and disk shelves located at the partner site are assigned to pool 1.

You should assign an equal number of shelves to each pool.

- a. On the first node at the remote site, systematically assign its local disk shelves to pool 0 and its remote disk shelves to pool 1:

```
disk assign -shelf local-switch-namesshelf-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.

#### 7.7.4. 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.

### 7.7.5. 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.

```
cluster_A::> network port show
```

| Node     | Port  | IPspace | Broadcast Domain | Link | MTU  | Speed (Mbps)<br>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.

## 7.7.6. Configure the clusters into a MetroCluster configuration

### Configure intercluster LIFs

#### Configuring intercluster LIFs on dedicated ports

You can configure intercluster LIFs on dedicated ports. Doing so typically increases the available bandwidth for replication traffic.

#### Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```
cluster01::> network port show
```

| Node         | Port | IPspace | Broadcast Domain | Link | MTU  | Speed (Mbps)<br>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
Vserver Group Failover
Targets

Cluster
cluster01 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
cluster01 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 | <pre>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</pre> |
| 9.5 and 9.4   | <pre>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</pre>                   |

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:

|                                                                          |
|--------------------------------------------------------------------------|
| <b>In ONTAP 9.6 and later:</b>                                           |
| <code>network interface show -service-policy default-intercluster</code> |
| <b>In ONTAP 9.5 and 9.4:</b>                                             |
| <code>network interface show -role intercluster</code>                   |

For complete command syntax, see the man page.

```
cluster01::> network interface show -service-policy default-intercluster
```

| Vserver   | Logical Interface | Status Admin/Oper | Network Address/Mask | Current Node | Current Port | Is 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:

|                                                                                    |
|------------------------------------------------------------------------------------|
| <b>In ONTAP 9.6 and later:</b>                                                     |
| <code>network interface show -service-policy default-intercluster -failover</code> |
| <b>In ONTAP 9.5 and 9.4:</b>                                                       |
| <code>network interface show -role intercluster -failover</code>                   |

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" on 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

```

## Configuring intercluster LIFs on shared data ports

You can configure intercluster LIFs on ports shared with the data network. Doing so reduces the number of ports you need for intercluster networking.

### Steps

1. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

```

cluster01::> network port show
Node Port IPspace Broadcast Domain Link MTU Speed (Mbps)
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 9.4:

```
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 9.4:**

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

For complete command syntax, see the man page.

```
cluster01::> network interface show -service-policy default-intercluster
```

| Vserver   | Logical Interface | Status Admin/Oper | Network Address/Mask | Current Node | Current Port | Is 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 9.4:**

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

For complete command syntax, see the man page.

The following example shows that the intercluster LIFs "cluster01\_icl01" and "cluster01\_icl02" on the "e0c" port will fail over to the "e0d" port.

```
cluster01::> network interface show -service-policy default-intercluster -failover
 Logical Home Failover Failover
Vserver Interface Node:Port Policy Group

cluster01
 cluster01_icl01 cluster01-01:e0c local-only 192.168.1.201/24
 Failover Targets: cluster01-01:e0c,
 cluster01-01:e0d
 cluster01_icl02 cluster01-02:e0c local-only 192.168.1.201/24
 Failover Targets: cluster01-02:e0c,
 cluster01-02:e0d
```

## Mirroring the root aggregates

You must mirror the root aggregates to provide data protection.

By default, the root aggregate is created as RAID-DP type aggregate. You can change the root aggregate from RAID-DP to RAID4 type aggregate. The following command modifies the root aggregate for RAID4 type aggregate:

```
storage aggregate modify -aggregate aggr_name -raidtype raid4
```



On non-ADP systems, the RAID type of the aggregate can be modified from the default RAID-DP to RAID4 before or after the aggregate is mirrored.

## Steps

1. Mirror the root aggregate:

```
storage aggregate mirror aggr_name
```

The following command mirrors the root aggregate for controller\_A\_1:

```
controller_A_1::> storage aggregate mirror aggr0_controller_A_1
```

This mirrors the aggregate, so it consists of a local plex and a remote plex located at the remote MetroCluster site.

2. Repeat the previous step for each node in the MetroCluster configuration.

## Implementing the MetroCluster configuration

You must run the `metrocluster configure -refresh true` command to start data protection on the nodes that you have added to a MetroCluster

configuration.

### About this task

You issue the `metrocluster configure -refresh true` command once, on one of the newly added nodes, to refresh the MetroCluster configuration. You do not need to issue the command on each of the sites or nodes.

The `metrocluster configure -refresh true` command automatically pairs the two nodes with the lowest system IDs in each of the two clusters as disaster recovery (DR) partners. In a four-node MetroCluster configuration, there are two DR partner pairs. The second DR pair is created from the two nodes with higher system IDs.

### Steps

1. Refresh the MetroCluster configuration:

- a. Enter advanced privilege mode:

```
set -privilege advanced
```

- b. Refresh the MetroCluster configuration on one of the new nodes:

```
metrocluster configure -refresh true
```

The following example shows the MetroCluster configuration refreshed on both DR groups:

```
controller_A_2:*> metrocluster configure -refresh true
[Job 726] Job succeeded: Configure is successful.
```

```
controller_A_4:*> metrocluster configure -refresh true
[Job 740] Job succeeded: Configure is successful.
```

- c. Return to admin privilege mode:

```
set -privilege admin
```

2. Verify the networking status on site A:

```
network port show
```

The following example shows the network port usage on a four-node MetroCluster configuration:



```

cluster_A::> network port show

```

| Node           | Port | IPspace | Broadcast Domain | Link | MTU  | Speed (Mbps)<br>Admin/Oper |
|----------------|------|---------|------------------|------|------|----------------------------|
| -----          |      |         |                  |      |      |                            |
| controller_A_1 |      |         |                  |      |      |                            |
|                | e0a  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0b  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0c  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0d  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0e  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0f  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0g  | Default | Default          | up   | 1500 | auto/1000                  |
| controller_A_2 |      |         |                  |      |      |                            |
|                | e0a  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0b  | Cluster | Cluster          | up   | 9000 | auto/1000                  |
|                | e0c  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0d  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0e  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0f  | Default | Default          | up   | 1500 | auto/1000                  |
|                | e0g  | Default | Default          | up   | 1500 | auto/1000                  |

14 entries were displayed.

3. Verify the MetroCluster configuration from both sites in the MetroCluster configuration:

a. Verify the configuration from site A:

`metrocluster show`

```

cluster_A::> metrocluster show

```

Configuration: IP fabric

| Cluster           | Entry Name          | State      |
|-------------------|---------------------|------------|
| -----             |                     |            |
| Local: cluster_A  | Configuration state | configured |
|                   | Mode                | normal     |
| Remote: cluster_B | Configuration state | configured |
|                   | Mode                | normal     |

b. Verify the configuration from site B:

`metrocluster show`

```

cluster_B::> metrocluster show

Configuration: IP fabric

Cluster Entry Name State

Local: cluster_B Configuration state configured
 Mode normal
Remote: cluster_A Configuration state configured
 Mode normal

```

## Creating a mirrored data aggregate on each node

You must create a mirrored data aggregate on each node in the DR group.

### About this task

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.
- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.

In systems using ADP, aggregates are created using partitions in which each drive is partitioned in to P1, P2 and P3 partitions.

- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.

### [Disk and aggregate management](#)

### Steps

1. Display a list of available spares:

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

2. Create the aggregate:

```
storage aggregate create -mirror true
```

If you are logged in to the cluster on the cluster management interface, you can create an aggregate on any node in the cluster. To ensure that the aggregate is created on a specific node, use the `-node` parameter or specify drives that are owned by that node.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)

- List of specific drives that are to be added to the aggregate
- Number of drives to include



In the minimum supported configuration, in which a limited number of drives are available, you must use the `force-small-aggregate` option to allow the creation of a three disk RAID-DP aggregate.

- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives that can be included in a RAID group
- Whether drives with different RPM are allowed

For more information about these options, see the `storage aggregate create` man page.

The following command creates a mirrored aggregate with 10 disks:

```
cluster_A::> storage aggregate create aggr1_node_A_1 -diskcount 10 -node node_A_1 -mirror
true
[Job 15] Job is queued: Create aggr1_node_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

3. Verify the RAID group and drives of your new aggregate:

```
storage aggregate show-status -aggregate aggregate-name
```

## Configuring 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

1. From the ONTAP cluster prompt, add the bridge to health monitoring:

- a. Add the bridge, using the command for your version of ONTAP:

| ONTAP version | Command                                                                                |
|---------------|----------------------------------------------------------------------------------------|
| 9.5 and later | <code>storage bridge add -address 0.0.0.0 -managed-by in-band -name bridge-name</code> |
| 9.4           | <code>storage bridge add -address bridge-ip-address -name bridge-name</code>           |

- b. Verify that the bridge has been added and is properly configured:

`storage bridge show`

It might take as long as 15 minutes to reflect all data because of the polling interval. The ONTAP health monitor can contact and monitor the bridge if the value in the "Status" column is "ok", and other information, such as the worldwide name (WWN), is displayed.

The following example shows that the FC-to-SAS bridges are configured:

```

controller_A_1::> storage bridge show

Bridge Symbolic Name Is Monitored Monitor Status Vendor Model
Bridge WWN

ATTO_10.10.20.10 atto01 true ok Atto FibreBridge 7500N
20000010867038c0
ATTO_10.10.20.11 atto02 true ok Atto FibreBridge 7500N
20000010867033c0
ATTO_10.10.20.12 atto03 true ok Atto FibreBridge 7500N
20000010867030c0
ATTO_10.10.20.13 atto04 true ok Atto FibreBridge 7500N
2000001086703b80

4 entries were displayed

controller_A_1::>

```

### 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
Last Checked On: 9/13/2018 20:41:37
```

| Component          | Result |
|--------------------|--------|
| nodes              | ok     |
| lifs               | ok     |
| config-replication | ok     |
| aggregates         | ok     |
| clusters           | ok     |
| connections        | ok     |

6 entries were displayed.

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:

```
ccluster_A:> metrocluster check aggregate show

Last Checked On: 8/5/2014 00:42:58

Node Aggregate Check Result

controller_A_1 controller_A_1_aggr0
 controller_A_1_aggr1
 controller_A_1_aggr2
controller_A_2 controller_A_2_aggr0
 controller_A_2_aggr1
 controller_A_2_aggr2

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.

## **Sending a custom AutoSupport message after adding nodes to the MetroCluster configuration**

You should issue an AutoSupport message to notify Lenovo 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](#).

## Chapter 8. Expanding 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 add four new nodes to the MetroCluster IP configuration as a third DR group. This creates a twelve-node MetroCluster configuration.

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.

### Before you begin

- The old and new nodes must be running the same version of ONTAP.
- This procedure describes the steps required to expand one four-node DR group. If you have an eight-node configuration (two DR groups), you must repeat the entire procedure for each DR group.
- You must ensure that the old and new platform models are supported for platform mixing.
- You must ensure that the old and new platform models are both supported by the IP switches.

#### [Lenovo Press](#)

- 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.

### 8.1. 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> |

### 8.2. Supported platform combinations when adding a second DR group

The following table shows the supported platform combinations for eight-node 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.
- The combinations in this table apply only to regular or permanent 8 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.

|                         |                              | 8-node DR group 2  |                   |                    |                              |
|-------------------------|------------------------------|--------------------|-------------------|--------------------|------------------------------|
|                         |                              | DM5000H<br>DM5000F | DG5000<br>DM5100F | DM7000H<br>DM7000F | DM7100H<br>DG7000<br>DM7100F |
| 8-node<br>DR group<br>1 | DM5000H<br>DM5000F           | Note 2             |                   |                    |                              |
|                         | DG5000<br>DM5100F            | Note 2             |                   |                    |                              |
|                         | DM7000H<br>DM7000F           |                    |                   | Note 1             |                              |
|                         | DM7100H<br>DG7000<br>DM7100F |                    |                   |                    | Note 1                       |
|                         |                              |                    |                   |                    |                              |

- **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.

### 8.3. Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify Lenovo 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.

## 8.4. Verifying the health of the MetroCluster configuration

You must verify the health and connectivity of the MetroCluster configuration prior to performing the transition

### Steps

1. Verify the operation of the MetroCluster configuration in ONTAP:

- a. Check whether the system is multipathed:

```
node run -node node-name sysconfig -a
```

- b. Check for any health alerts on both clusters:

```
system health alert show
```

- c. Confirm the MetroCluster configuration and that the operational mode is normal:

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

2. Verify that the cluster is healthy:

```
cluster show
```

```

cluster_A::> cluster show
Node Health Eligibility

node_A_1 true true
node_A_2 true true

cluster_A::>

```

3. Verify that all cluster ports are up:

`network port show -ipSpace Cluster`

```

cluster_A::> network port show -ipSpace Cluster

Node: node_A_1-old

Port IPspace Broadcast Domain Link MTU Speed(Mbps) Health
Admin/Oper Status

e0a Cluster Cluster up 9000 auto/10000 healthy
e0b Cluster Cluster up 9000 auto/10000 healthy

Node: node_A_2-old

Port IPspace Broadcast Domain Link MTU Speed(Mbps) Health
Admin/Oper Status

e0a Cluster Cluster up 9000 auto/10000 healthy
e0b Cluster Cluster up 9000 auto/10000 healthy

4 entries were displayed.

cluster_A::>

```

4. Verify that all cluster LIFs are up and operational:

`network interface show -vserver Cluster`

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

```
cluster_A::> network interface show -vserver cluster
```

| Vserver | Logical Interface  | Status Admin/Oper | Network Address/Mask | Current Node | Current Port | Is 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
```

| Vserver | Logical 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::>
```

### 8.5. Removing the configuration from monitoring applications

If the existing configuration is monitored with the MetroCluster Tiebreaker software, the ONTAP Mediator or other third-party applications (for example, ClusterLion) that can initiate a switchover, you must remove the MetroCluster configuration from the monitoring software prior to upgrade.

**Steps**

1. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software

that can initiate switchover.

| If you are using... | Use this procedure...                                                                                                      |
|---------------------|----------------------------------------------------------------------------------------------------------------------------|
| Tiebreaker          | <a href="#">Removing MetroCluster Configurations.</a>                                                                      |
| Mediator            | Issue the following command from the ONTAP prompt:<br><br><code>metrocluster configuration-settings mediator remove</code> |

2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.

Refer to the documentation for the application.

## 8.6. 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 the *MetroCluster installation and configuration*.

#### [Cabling the IP switches](#)

3. Configure the MetroCluster IP nodes using the following sections of the *MetroCluster installation and configuration*.
  - a. [Gathering required information](#)

- b. [Restoring system defaults on a controller module](#)
  - c. [Verifying the ha-config state of components](#)
  - d. [Manually assigning 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.

## 8.7. 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

Last Checked On: 4/7/2019 21:15:05

Component Result

nodes ok
lifs ok
config-replication ok
aggregates warning
clusters ok
connections not-applicable
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:

- [Resetting the Broadcom IP switch to factory defaults](#)
- [Resetting the Cisco IP 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

- [Downloading and installing the Broadcom IP RCF files](#)
- [Downloading and installing the Cisco IP RCF files](#)



If you have an L2 shared or L3 network configuration, you might need to adjust the ISL ports on the intermediate/customer switches. The switchport mode might change from 'access' to 'trunk' mode. Only proceed to upgrade the second switch pair (A\_2, B\_2) if the network connectivity between switches A\_1 and B\_1 is fully operational and the network is healthy.

## 8.8. Joining 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.
```

```
IMPORTANT: Go to https://www.lenovo.com/registration to register your device.
```

```
By enabling AutoSupport, you agree that event messages and periodic system
operational data will be transmitted to Lenovo. This data may include email
addresses you designate to receive AutoSupport notifications, system usernames,
and device identifiers such as MAC addresses, and Node and SVM names.
```

```
Enabling AutoSupport also enables AutoSupport OnDemand. This feature allows the
Support team to direct your system to transmit additional technical support
data, including logs and core dump files. AutoSupport OnDemand feature can be
disabled using the Command Line Interface (CLI) by entering:
"system node autosupport modify -node * -ondemand-state disable", in advanced
mode.
```

```
AutoSupport and AutoSupport OnDemand data: (a) is transmitted to a U.S. based
data center operated by NetApp, Lenovo's service provider; (b) will be used
to enhance the technical support and services provided to you; and (c) Support
analytics to help maintain the health and performance of your system.
To understand how Lenovo handles data, please see our Privacy Statements at
www.lenovo.com/privacy.
```

```
Enabling AutoSupport can significantly speed problem determination
and resolution, should a problem occur on your system.
```

```
To enable AutoSupport data transmission, type "yes".
Leave this field blank to keep data transmission disabled {yes, no} [no]:
```

```
To continue with the cluster setup wizard, type "yes" {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.

## 8.9. Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates

You must create cluster peering LIFs, create the MetroCluster interfaces on the new MetroCluster IP nodes.

### About this task

The home port used in the examples are platform-specific. You should use the appropriate home

port specific to MetroCluster IP node platform.

## Steps

1. On the new MetroCluster IP nodes, configure the intercluster LIFs 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
 cluster_A node_A_2-old node_B_2-old
 cluster_B node_B_1-old node_A_1-old
 cluster_B node_B_2-old node_A_2-old
2 cluster_A node_A_1-new node_B_1-new
 cluster_A node_A_2-new node_B_2-new
 cluster_B node_B_1-new node_A_1-new
 cluster_B 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:

```
metrocluster configuration-settings interface create -cluster-name
```

**Notes:**

- Starting with ONTAP 9.8, certain platforms use a VLAN for the MetroCluster IP interface. By default, each of the two ports uses a different VLAN: 10 and 20. You can also specify a different (non-default) VLAN higher than 100 (between 101–4095) using the `-vlan-id` parameter in the `metrocluster configuration-settings interface create` command.
- Starting with ONTAP 9.9.1, if you are using a layer 3 configuration, you must also specify the `-gateway` parameter when creating MetroCluster IP interfaces. Refer to [Considerations for layer 3 wide-area networks](#).

The following platform models can be added to the existing MetroCluster configuration if the VLANs used are 10/20 or greater than 100. If any other VLANs are used, then these platforms cannot be added to the existing configuration as the MetroCluster interface cannot be configured. If you are using any other platform, the VLAN configuration is not relevant as this is not required in ONTAP.

| AFA platforms                                                                                           | Hybrid platforms                                                           |
|---------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>AFA DM5000F</li> <li>AFA DM5100F</li> <li>AFA DM7100F</li> </ul> | <ul style="list-style-type: none"> <li>DM5000H</li> <li>DM7100H</li> </ul> |



You can configure the MetroCluster IP interfaces from either cluster.

```
cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_1-new -home-port e1a -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 e1b -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 e1a -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 e1b -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 e1a -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 e1b -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 e1a -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 e1b -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 | Group | Cluster   | Node           | Network Address | Netmask       | Gateway       | Config 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: e1a  | 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: e1a  | 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: e1a  | 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 Group | Cluster   | Node         | Source Network Address         | Destination Network Address | Partner      | Type | Config    | State |
|----------|-----------|--------------|--------------------------------|-----------------------------|--------------|------|-----------|-------|
| 1        | cluster_A | node_A_1-old | Home Port: e1a<br>172.17.28.10 | 172.17.28.11                | HA Partner   |      | completed |       |
|          |           |              | Home Port: e1a<br>172.17.28.10 | 172.17.28.12                | DR Partner   |      | completed |       |
|          |           |              | Home Port: e1a<br>172.17.28.10 | 172.17.28.13                | DR Auxiliary |      | completed |       |
|          |           |              | Home Port: e1b<br>172.17.29.10 | 172.17.29.11                | HA Partner   |      | completed |       |
|          |           |              | Home Port: e1b<br>172.17.29.10 | 172.17.29.12                | DR Partner   |      | completed |       |
|          |           |              | Home Port: e1b<br>172.17.29.10 | 172.17.29.13                | DR Auxiliary |      | completed |       |
|          |           | node_A_2-old | Home Port: e1a<br>172.17.28.11 | 172.17.28.10                | HA Partner   |      | completed |       |
|          |           |              | Home Port: e1a<br>172.17.28.11 | 172.17.28.13                | DR Partner   |      | completed |       |
|          |           |              | Home Port: e1a<br>172.17.28.11 | 172.17.28.12                | DR Auxiliary |      | completed |       |
|          |           |              | Home Port: e1b<br>172.17.29.11 | 172.17.29.10                | HA Partner   |      | completed |       |
|          |           |              | Home Port: e1b<br>172.17.29.11 | 172.17.29.13                | DR Partner   |      | completed |       |
|          |           |              | Home Port: e1b<br>172.17.29.11 | 172.17.29.12                | DR Auxiliary |      | completed |       |

| DR Group | Cluster   | Node         | Source Network Address         | Destination Network Address | Partner    | Type | Config    | State |
|----------|-----------|--------------|--------------------------------|-----------------------------|------------|------|-----------|-------|
| 1        | cluster_B | node_B_2-old | Home Port: e1a<br>172.17.28.13 | 172.17.28.12                | HA Partner |      | completed |       |
|          |           |              | Home Port: e1a<br>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 | completed    |
| DR    |                | Source         | Destination     |                 |              |              |
| Group | Cluster        | Node           | Network Address | Network Address | Partner Type | Config State |
| ----- |                |                |                 |                 |              |              |
| 2     | cluster_A      |                |                 |                 |              |              |
|       | node_A_1-new** |                |                 |                 |              |              |
|       |                | Home Port: e1a | 172.17.26.10    | 172.17.26.11    | HA Partner   | completed    |
|       |                | Home Port: e1a | 172.17.26.10    | 172.17.26.12    | DR Partner   | completed    |
|       |                | Home Port: e1a | 172.17.26.10    | 172.17.26.13    | DR Auxiliary | completed    |
|       |                | Home Port: e1b | 172.17.27.10    | 172.17.27.11    | HA Partner   | completed    |
|       |                | Home Port: e1b | 172.17.27.10    | 172.17.27.12    | DR Partner   | completed    |
|       |                | Home Port: e1b | 172.17.27.10    | 172.17.27.13    | DR Auxiliary | completed    |
|       | node_A_2-new   |                |                 |                 |              |              |
|       |                | Home Port: e1a | 172.17.26.11    | 172.17.26.10    | HA Partner   | completed    |
|       |                | Home Port: e1a | 172.17.26.11    | 172.17.26.13    | DR Partner   | completed    |
|       |                | Home Port: e1a | 172.17.26.11    | 172.17.26.12    | DR Auxiliary | completed    |
|       |                | Home Port: e1b | 172.17.27.11    | 172.17.27.10    | HA Partner   | completed    |
|       |                | Home Port: e1b | 172.17.27.11    | 172.17.27.13    | DR Partner   | completed    |
|       |                | Home Port: e1b | 172.17.27.11    | 172.17.27.12    | DR Auxiliary | completed    |
| DR    |                | 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_rha1_a1 node_a_1
2.20.3 546.9GB 20 3 SAS aggregate aggr0_rha1_a2 node_a_2
2.20.5 546.9GB 20 5 SAS aggregate rha1_a1_aggr1 node_a_1
2.20.6 546.9GB 20 6 SAS aggregate rha1_a1_aggr1 node_a_1
2.20.7 546.9GB 20 7 SAS aggregate rha1_a2_aggr1 node_a_2
2.20.10 546.9GB 20 10 SAS aggregate rha1_a1_aggr1 node_a_1
...
43 entries were displayed.

cluster_A::>

```

10. Mirror the root aggregates:

`storage aggregate mirror -aggregate aggr0_node_A_1-new`



You must complete this step on each MetroCluster IP node.

```

cluster_A::> aggr mirror -aggregate aggr0_node_A_1-new

Info: Disks would be added to aggregate "aggr0_node_A_1-new" on node "node_A_1-new"
 in the following manner:

 Second Plex

 RAID Group rg0, 3 disks (block checksum, raid_dp)

 Position Disk Type Usable Physical

 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,<br>mirrored,<br>normal |
| aggr0_node_A_2-old | 349.0GB | 16.84GB   | 95%   | online | 1     | node_A_2-old | raid_dp,<br>mirrored,<br>normal |
| aggr0_node_A_1-new | 467.6GB | 22.63GB   | 95%   | online | 1     | node_A_1-new | raid_dp,<br>mirrored,<br>normal |
| aggr0_node_A_2-new | 467.6GB | 22.62GB   | 95%   | online | 1     | node_A_2-new | raid_dp,<br>mirrored,<br>normal |
| aggr_data_a1       | 1.02TB  | 1.01TB    | 1%    | online | 1     | node_A_1-old | raid_dp,<br>mirrored,<br>normal |
| aggr_data_a2       | 1.02TB  | 1.01TB    | 1%    | online | 1     | node_A_2-old | raid_dp,<br>mirrored,           |

## 8.10. 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 one of the new nodes:

```
metrocluster configure
```

The following example shows the MetroCluster configuration refreshed on both DR groups:

```
cluster_A::*> metrocluster configure -refresh true
[Job 726] Job succeeded: Configure is successful.
```

c. 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 support that one site of the MetroCluster has a single mirrored data aggregate and the other site has more than one mirrored data aggregate.

The following example shows the creation of an aggregate on node\_A\_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)

| Position | Disk    | Type | Usable<br>Size | Physical<br>Size |
|----------|---------|------|----------------|------------------|
| dparity  | 5.10.15 | SAS  | -              | -                |
| parity   | 5.10.16 | SAS  | -              | -                |
| data     | 5.10.17 | SAS  | 546.9GB        | 547.1GB          |
| data     | 5.10.18 | SAS  | 546.9GB        | 558.9GB          |
| data     | 5.10.19 | SAS  | 546.9GB        | 558.9GB          |

Second Plex

RAID Group rg0, 5 disks (block checksum, raid\_dp)

| Position | Disk    | Type | Usable<br>Size | Physical<br>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 from the old nodes to the new nodes in advanced privilege.
  - a. Display the volumes to identify the MDV volumes:



If you have a single mirrored data aggregate per site then move both the MDV volumes to this single aggregate. If you have two or more mirrored data aggregates, then move each MDV volume to a different aggregate.

The following example shows the MDV volumes in the `volume show` output:

```

cluster_A:::> volume show
Vserver Volume Aggregate State Type Size Available Used%

...
cluster_A MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_A
 agr_b1 - RW - - -
cluster_A MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_B
 agr_b2 - RW - - -
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_A
 agr_a1 online RW 10GB 9.50GB 0%
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
 agr_a2 online RW 10GB 9.50GB 0%
...
11 entries were displayed.mple

```

- b. Set the advanced privilege level:

```
set -privilege advanced
```

- c. Move the MDV volumes, one at a time:

```
volume move start -volume mdv-volume -destination-aggregate aggr-on-new-node
-vserver vsERVER-name
```

The following example shows the command and output for moving "MDV\_CRS\_d6b0b313ff5611e9837100a098544e51\_A" to aggregate "data\_a3" on "node\_A\_3".

```
cluster_A::*> vol move start -volume MDV_CRS_d6b0b313ff5611e9837100a098544e51_A -destination
-aggregate data_a3 -vserver cluster_A
```

```
Warning: You are about to modify the system volume
"MDV_CRS_d6b0b313ff5611e9837100a098544e51_A". This might cause severe
performance or stability problems. Do not proceed unless directed to
do so by support. Do you want to proceed? {y|n}: y
[Job 494] Job is queued: Move "MDV_CRS_d6b0b313ff5611e9837100a098544e51_A" in Vserver
"cluster_A" to aggregate "data_a3". Use the "volume move show -vserver cluster_A -volume
MDV_CRS_d6b0b313ff5611e9837100a098544e51_A" command to view the status of this operation.
```

- d. Use the volume show command to check that the MDV volume has been successfully moved:

```
volume show mdv-name
```

The following output shows that the MDV volume has been successfully moved.

```
cluster_A::*> vol show MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
Vserver Volume Aggregate State Type Size Available Used%

cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
 aggr_a2 online RW 10GB 9.50GB 0%
```

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

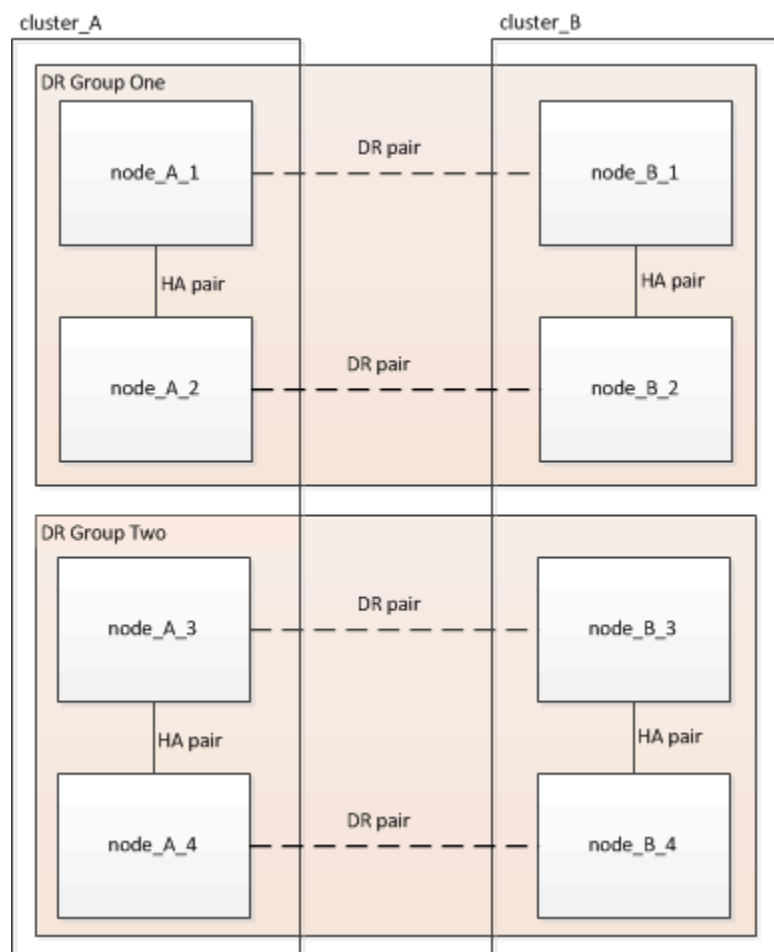
node_A_1-old false
node_A_2-old false
node_A_3-new true
node_A_4-new false
4 entries were displayed.
```

## Chapter 9. Removing a Disaster Recovery group


Beginning with ONTAP 9.8, you can remove a DR group from an eight-node MetroCluster configuration to create a four-node MetroCluster configuration.

This procedure is supported on ONTAP 9.8 and later.

An eight-node configuration includes eight-nodes organized as two four-node DR groups.



By removing a DR Group, four nodes remain in the configuration.



## 9.1. Removing the DR group nodes from each cluster

- You must perform this step on both clusters.
- The `metrocluster remove-dr-group` command is supported only on ONTAP 9.8 and later.
  1. Prepare for the removal of the DR group, if you haven't already.
    - a. Move all data volumes to another DR group.
    - b. Move all MDV\_CRS metadata volumes to another DR group. Follow the steps in the following procedure: [Moving a metadata volume in MetroCluster configurations](#)
    - c. Delete all MDV\_aud metadata volumes that might exist in the DR group to be removed.
    - d. Delete all data aggregates in the DR group to be removed as shown in the following example:

```
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.

- e. Move the data LIFs offline. `network interface modify -vserver svm-name -lif data-lif -status-admin down`
- f. Migrate all data LIFs to home nodes in another DR group.
 

```
network interface show -home-node old_node
```

```
network interface modify -vserver svm-name -lif data-lif -home-node new_node
-home-port port-id
```

- g. Move the data LIFs back online. `network interface modify -vserver svm-name -lif data-lif -status-admin up`
- 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.

- 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.

```

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. If 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 the clusters.

- a. Disconnect the connections:

```
metrocluster configuration-settings connection disconnect dr-group-id
```

- b. Delete the MetroCluster interfaces on the nodes of the old DR group:

```
metrocluster configuration-settings interface delete
```

- c. Delete the old DR group's configuration.

```
metrocluster configuration-settings dr-group delete
```

5. Unjoin the nodes in the old DR group.

You must perform this step on each cluster.

- a. Set the advanced privilege level:

```
set -privilege advanced
```

- b. Disable the storage failover:

```
storage failover modify -node node-name -enable false
```

- c. Unjoin the node:

```
cluster unjoin -node node-name
```

Repeat this step for the other local node in the old DR group.

- d. Set the admin privilege level:

```
set -privilege admin
```

6. Re-enable cluster HA in the new DR group:

```
cluster ha modify -configured true
```

You must perform this step on each cluster.


7. Halt, power down, and remove the old controller modules and storage shelves.

## Chapter 10. Where to find additional information

You can learn more about MetroCluster configuration and operation.

### 10.1. MetroCluster and miscellaneous information

| Information                                                                 | Subject                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">ONTAP 9 Documentation Center</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>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <a href="#">MetroCluster management and disaster recovery</a>               | <ul style="list-style-type: none"><li>• Understanding the MetroCluster configuration</li><li>• Switchover, healing and switchback</li><li>• Disaster recovery</li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <a href="#">Maintain MetroCluster Components</a>                            | <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 four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.</li></ul> |

|                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <a href="#">MetroCluster Upgrade, Transition, and Expansion</a>                                                                                                                                                                                                                                                                                                                                             | <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> |
| <a href="#">MetroCluster Tiebreaker Software installation and configuration</a>                                                                                                                                                                                                                                                                                                                             | <ul style="list-style-type: none"> <li>• Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software</li> </ul>                                                                                                                                                |
| <a href="#">Lenovo ThinkSystem Storage Documentation Center</a> <div data-bbox="240 772 310 842" style="float: left; margin-right: 10px;">  </div> <div data-bbox="386 737 781 884" style="clear: both;"> <p>The standard storage shelf maintenance procedures can be used with MetroCluster IP configurations.</p> </div> | <ul style="list-style-type: none"> <li>• Hot-adding a disk shelf</li> <li>• Hot-removing a disk shelf</li> </ul>                                                                                                                                                                       |
| <a href="#">ONTAP concepts</a>                                                                                                                                                                                                                                                                                                                                                                              | <ul style="list-style-type: none"> <li>• How mirrored aggregates work</li> </ul>                                                                                                                                                                                                       |



# Chapter 11. Appendix

## 11.1. Contacting support

You can contact Support to obtain help for your issue.

You can receive hardware service through a Lenovo Authorized Service Provider. To locate a service provider authorized by Lenovo to provide warranty service, go to <https://datacentersupport.lenovo.com/serviceprovider> and use filter searching for different countries. For Lenovo support telephone numbers, see <https://datacentersupport.lenovo.com/supportphonenumber> for your region support details.

## 11.2. Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area.

Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document is not an offer and does not provide a license under any patents or patent applications. You can send inquiries in writing to the following:

Lenovo (United States), Inc.  
8001 Development Drive  
Morrisville, NC 27560  
U.S.A.  
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information

contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

### **11.3. Trademarks**

LENOVO, LENOVO logo, and THINKSYSTEM are trademarks of Lenovo. All other trademarks are the property of their respective owners. © 2024 Lenovo.

**Lenovo**