A photograph of Fujitsu ETERNUS storage hardware. The top part shows a close-up of a metal grille with a diamond-shaped pattern. Below it, a rack-mounted unit is visible with the "ETERNUS" logo on its front panel. The background is a light gray gradient.

FUJITSU Storage ETERNUS AX/HX Series MetroCluster[®] Upgrade, Transition, and Expansion Guide

Contents

Deciding which procedure to use.....	5
Choosing an upgrade or refresh method.....	6
Upgrading controllers in a MetroCluster IP configuration using switchover and switchback (ONTAP 9.8 and later).....	7
Workflow for upgrading controllers in an MetroCluster IP configuration.....	7
Preparing for the upgrade.....	8
Updating the MetroCluster switch RCF files before upgrading controllers.....	8
Mapping ports from the old nodes to the new nodes.....	9
Netbooting the new controllers.....	10
Clearing the configuration on a controller module.....	12
Verifying MetroCluster health before site upgrade.....	12
Gathering information before the upgrade.....	13
Removing Mediator or Tiebreaker monitoring.....	14
Sending a custom AutoSupport message prior to maintenance.....	14
Switching over the MetroCluster configuration.....	15
Removing interface configurations and uninstalling the old controllers.....	15
Updating the switch RCFs to accommodate the new platforms.....	16
Configuring the new controllers.....	17
Setting up the new controllers.....	17
Restoring the HBA configuration.....	17
Setting the HA state on the new controllers and chassis.....	18
Setting the MetroCluster IP bootarg variables.....	18
Reassigning root aggregate disks.....	20
Booting up the new controllers.....	21
Verifying and restoring LIF configuration.....	22
Switching back the MetroCluster configuration.....	22
Checking the health of the MetroCluster configuration.....	24
Upgrading the nodes on cluster_A.....	24
Restoring Tiebreaker or Mediator monitoring.....	24
Sending a custom AutoSupport message after maintenance.....	24
Refreshing a four-node MetroCluster FC configuration.....	26
Refreshing a four-node MetroCluster IP configuration (ONTAP 9.8 and later).....	29
Adding new nodes to the MetroCluster IP configuration.....	31
Example naming in this procedure.....	32
Sending a custom AutoSupport message prior to maintenance.....	32
Verifying the health of the MetroCluster configuration.....	32
Removing the configuration from monitoring applications.....	34
Preparing the new controller modules.....	34
Joining the new nodes to the clusters.....	35
Configuring intercluster LIFs, creating the MetroCluster interfaces, and mirroring root aggregates.....	36
Finalizing the addition of the new nodes.....	40
Transitioning nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later).....	44
Preparing for transition from a MetroCluster FC to a MetroCluster IP configuration.....	44
Supported platforms for nondisruptive transition.....	44
Requirements for nondisruptive FC-to-IP transition.....	44
How transition impacts the MetroCluster hardware components.....	44
Workflow for nondisruptive MetroCluster transition.....	45
Considerations for IP switches.....	47

Switchover, healing, and switchback operations during transition.....	47
Alert messages and tool support during transition.....	48
Example naming in this procedure.....	48
Transitioning from MetroCluster FC to MetroCluster IP configurations.....	49
Verifying the health of the MetroCluster configuration.....	49
Removing the existing configuration from the Tiebreaker or other monitoring software.....	50
Generating and applying RCFs to the new IP switches.....	51
Moving the local cluster connections.....	51
Preparing the MetroCluster IP controllers.....	56
Configure the MetroCluster for transition.....	56
Moving the data to the new drive shelves.....	65
Removing the MetroCluster FC controllers.....	65
Completing the transition.....	68
Sending a custom AutoSupport message after maintenance.....	69
Restoring Tiebreaker or Mediator monitoring.....	69
Moving SAN FC hosts from MetroCluster FC to MetroCluster IP nodes.....	70
Moving Linux iSCSI hosts from MetroCluster FC to MetroCluster IP nodes..	74
Setting up new iSCSI connections.....	74
Adding the MetroCluster IP nodes as reporting nodes.....	76
Removing reporting nodes and rescanning paths.....	78
Expanding a two-node MetroCluster FC configuration to a four-node configuration.....	80
Verifying the state of the MetroCluster configuration.....	82
Sending a custom AutoSupport message before adding nodes to the MetroCluster configuration.....	83
Zoning for the new controller ports when adding a controller module in a fabric-attached MetroCluster configuration.....	83
Adding a new controller module to each cluster.....	83
Preparing for the upgrade.....	85
Clearing the configuration on a controller module.....	85
Preparing cluster ports on an existing controller module.....	86
Preparing the netboot server to download the image.....	87
Setting the HA mode on the existing controller module.....	88
Shutting down the existing controller module.....	88
Installing and cabling the new controller module.....	88
Powering up both controller modules and displaying the LOADER prompt.....	90
Changing the ha-config setting on the existing and new controller modules.....	91
Setting the partner system ID for both controller modules.....	91
Booting the existing controller module.....	91
Assigning disks to the new controller module.....	91
Netbooting and setting up ONTAP on the new controller module.....	92
Mirroring the root aggregate on the new controller.....	94
Configuring intercluster LIFs.....	94
Creating a mirrored data aggregate on each node.....	98
Installing licenses for the new controller module.....	98
Creating unmirrored data aggregates.....	99
Installing the firmware after adding a controller module.....	100
Refreshing the MetroCluster configuration with new controllers.....	100
Enabling storage failover on both controller modules and enabling cluster HA.....	101
Restarting the SVMs.....	102
Expanding a four-node MetroCluster FC configuration to an eight-node configuration.....	103
Determining the new cabling layout.....	104
Racking the new equipment.....	105
Verifying the health of the MetroCluster configuration.....	105
Checking for MetroCluster configuration errors with Config Advisor.....	106
Sending a custom AutoSupport message prior to adding nodes to the MetroCluster configuration.....	106
Recabling and zoning a switch fabric for the new nodes.....	106

Disconnecting the existing DR group from the fabric.....	106
Applying the RCF files and recabling the switches.....	107
Configuring ONTAP on the new controllers.....	107
Restoring system defaults on a previously used controller module.....	107
Assigning disk ownership in ETERNUS AX systems.....	109
Assigning disk ownership in non-ETERNUS AX systems.....	110
Verifying the ha-config state of components.....	111
Booting the new controllers and joining them to the cluster.....	112
Configuring the clusters into a MetroCluster configuration.....	113
Checking for MetroCluster configuration errors with Config Advisor.....	122
Sending a custom AutoSupport message after to adding nodes to the MetroCluster configuration....	122
Verifying switchover, healing, and switchback.....	123
Removing a Disaster Recovery group.....	124
Where to find additional information.....	127
Copyright and trademark.....	128
Copyright.....	128
Trademark.....	128
How to send comments about documentation and receive update notifications.....	128

Deciding which procedure to use

You must understand the differences between a MetroCluster upgrade, expansion, or transition and choose the procedure that matches your goal.

Upgrade

In an upgrade procedure, you are replacing the controller modules with a new model of controller module.

- The old controller modules are retired.
- The storage is not upgraded.
- The storage switch infrastructure technology is not changed.

It remains as a MetroCluster IP.

[Choosing an upgrade or refresh method](#) on page 6

Refresh

In a refresh procedure, you are replacing the controller modules with a new model of controller module and also replacing the storage shelves.

- The old controller modules and storage shelves are retired.
- The storage switch infrastructure technology is not changed, unless new switches are required to accommodate the new platform models.

The configuration remains as its original type:

- MetroCluster IP

[Choosing an upgrade or refresh method](#) on page 6

Transition

In a transition procedure, the backend switch infrastructure is changed from FC switches or connections to IP switches, and the platform models are replaced.

- New controller modules are added to the configuration.
- The original controller modules are retired after the procedure.
- If the original configuration used cluster interconnect switches, they can be reused, depending on the switch and platform models.
- The storage can be reused or replaced, depending on the platform models.

If the storage shelves are not reused, data is moved from the old shelves to the new shelves.

[Transitioning nondisruptively from a MetroCluster FC to a MetroCluster IP configuration \(ONTAP 9.8 and later\)](#) on page 44

Expansion

In an expansion procedure, additional nodes and storage are added to the MetroCluster FC configuration.

- A two-node MetroCluster FC configuration can be expanded to a four-node configuration.
- A four-node MetroCluster FC configuration can be expanded to an eight-node configuration consisting of two DR groups.

[Expanding a two-node MetroCluster FC configuration to a four-node configuration](#) on page 80

[Expanding a four-node MetroCluster FC configuration to an eight-node configuration](#) on page 103

Choosing an upgrade or refresh method

The upgrade or refresh procedure you use depends on the platform model, scope of the upgrade, and type of MetroCluster configuration.

There are different types of upgrade and refresh procedures.

- 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 storage platform and the storage shelves.

In the refresh procedures, new nodes 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 nodes are then retired.

Type of upgrade or refresh		MetroCluster type	First ONTAP version support	Procedure
<ul style="list-style-type: none"> • Scope: Platform (controller modules) only • Method: Switchover/switchback 		FC	9.8	Upgrading controllers in an MetroCluster FC configuration using switchover and switchback (ONTAP 9.8 and later)
		IP	9.8	Upgrading controllers in a MetroCluster IP configuration using switchover and switchback (ONTAP 9.8 and later) on page 7
<ul style="list-style-type: none"> • Scope: Platform (controller modules) and storage shelves • Method: Switchover/switchback 		FC	9.7 and later	Refreshing a four-node MetroCluster FC configuration on page 26
		IP	9.8	Refreshing a four-node MetroCluster IP configuration (ONTAP 9.8 and later) on page 29
<ul style="list-style-type: none"> • Scope: Platform (controller modules) • Method: Aggregate relocation (ARL) 	Using system controller replace commands Can be used with ONTAP 9.7 and later	FC	9.7 and later	
	Using manual ARL commands	FC	9.x	

Upgrading controllers in a MetroCluster IP configuration using switchover and switchback (ONTAP 9.8 and later)

Starting 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.

Before you begin

- The platforms must be running ONTAP 9.8 or later.
- This procedure applies to controller modules in a MetroCluster IP configuration.
- The supported upgrade path depends on the original platform model.

Platform models with internal shelves are not supported.

Old platform model	New platform model
<ul style="list-style-type: none">• AX3100	<ul style="list-style-type: none">• AX4100
<ul style="list-style-type: none">• HX6100	<ul style="list-style-type: none">• FAS9000• HX6100

- 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.

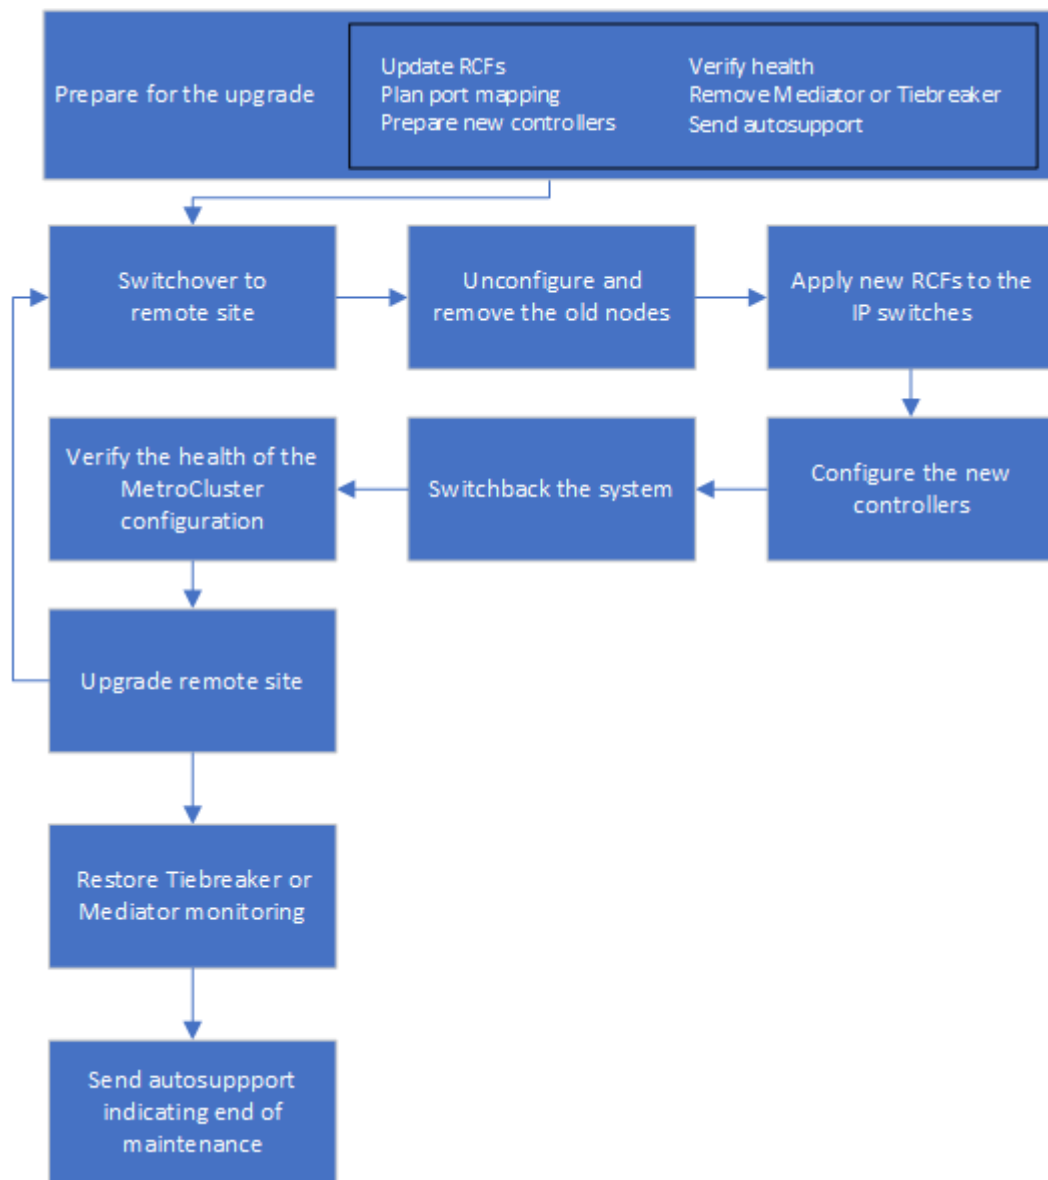
- The new platform must be a different model than the original platform.
- The IP switches must be running a supported firmware version.
- If the new platform has fewer slots than the original system, or if it has fewer or different types of ports, you might need to add an adapter to the new system.

About this task

- 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

Workflow for upgrading controllers in an MetroCluster IP configuration

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



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.

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.

Note: 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"> AX4100 	<ul style="list-style-type: none"> 10 20 Any value in the range 101 to 4096 inclusive.

- 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.

Procedure

1. 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 Guide](#).

- 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 [MetroCluster IP Installation and Configuration Guide](#).

- Downloading and installing the Broadcom RCF files
- Downloading and installing the Cisco IP RCF files

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.

Cluster interconnect physical ports		
Old controller	New controller	Required action
e0c, e0d	e0a,e0b,e0c,e0d	e0c and e0d are matching ports. You do not have to change the configuration, but after upgrade you can spread your cluster LIFs across the available cluster ports.

Procedure

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.
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						

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.

Procedure

Netboot the new controllers:

- a) Download the appropriate ONTAP software from the DVD included in the Product and store the *ontap-version_image.tgz* file on a web-accessible directory.
- b) Change to the web-accessible directory and verify that the files you need are available.

If the platform model is...	Then...
AX/HX 6000 series systems	<p>Extract the contents of the <code>ontap-version_image.tgz</code> file to the target directory: <code>tar -zxvf ontap-version_image.tgz</code></p> <p>Note: If you are extracting the contents on Windows, use 7-Zip or WinRAR to extract the netboot image. Your directory listing should contain a netboot folder with a kernel file: <code>netboot/kernel</code></p> <p>Your directory listing should contain a netboot folder with a kernel file:</p> <pre>netboot/kernel</pre>
All other systems	<p>Your directory listing should contain a netboot folder with a kernel file:</p> <pre>ontap-version_image.tgz</pre> <p>You do not need to extract the <code>ontap-version_image.tgz</code> file.</p>

c) At the LOADER prompt, configure the netboot connection for a management LIF:

If IP addressing is...	Then...
DHCP	<p>Configure the automatic connection:</p> <pre>ifconfig e0M -auto</pre>
Static	<p>Configure the manual connection:</p> <pre>ifconfig e0M -addr=<i>ip_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i></pre>

d) Perform the netboot.

If the platform model is...	Then...
AX/HX 6000 series systems	<pre>netboot http://web_server_ip/path_to_web-accessible_directory/netboot/kernel</pre>
All other systems	<pre>netboot http://web_server_ip/path_to_web-accessible_directory/ontap-version_image.tgz</pre>

e) 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.

f) 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.

g) 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
```

h) 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}
```

Clearing the configuration on a controller module

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

Procedure

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: `saveenvbye`
4. At the LOADER prompt, launch the boot menu: `boot_ontap menu`
5. At the boot menu prompt, clear the configuration: `wipeconfig`
Respond `yes` to the confirmation prompt.
The node reboots and the boot menu is displayed again.
6. At the boot menu, select option **5** to boot the system into Maintenance mode.
Respond `yes` to the confirmation prompt.

Verifying MetroCluster health before site upgrade

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

Procedure

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 timezone 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 timezone.
2. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.
 - a) Confirm the MetroCluster configuration and that the operational mode is normal: `metrocluster show`
 - b) Confirm that all expected nodes are shown: `metrocluster node show`
 - c) Issue the following command: `metrocluster check run`
 - d) Display the results of the MetroCluster check: `metrocluster check show`
3. Check the MetroCluster cabling with the Config Advisor tool.
 - a) Download and run Config Advisor.
 - b) After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

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.

Procedure

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 -nodenode-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 -nodenode-namesysconfig -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

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

You should gather the output of the following commands:

- `fcp adapter show -instance`
- `fcp interface show -instance`
- `iscsi interface show`

- `ucadmin show`
- 5. If the root volume is encrypted, collect and save the passphrase used for key-manager: `security key-manager backup show`
- 6. If the MetroCluster nodes are using encryption for volumes or aggregates, copy information about the keys and passphrases.
 - 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 -instancesecurity key-manager key query`
- 7. Gather the system IDs of the existing nodes: `metrocluster node show -fields node-systemid,ha-partner-systemid,dr-partner-systemid,dr-auxiliary-systemid`
The following output shows the reassigned drives.

```

::> metrocluster node show -fields node-systemid,ha-partner-systemid,dr-partner-systemid,dr-auxiliary-systemid
dr-group-id cluster      node      node-systemid ha-partner-systemid dr-partner-systemid dr-auxiliary-systemid
-----
1            cluster_A node_A_1  537403324   537403323         537403321         537403322
1            cluster_A node_A_2  537403323   537403324         537403322         537403321
1            cluster_B node_B_1  537403322   537403321         537403323         537403324
1            cluster_B node_B_2  537403321   537403322         537403324         537403323
4 entries were displayed.
    
```

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.

Procedure

1. Collect the output for the following command: `storage iscsi-initiator show`
2. Remove the existing MetroCluster configuration from Tiebreaker, Mediator, or other software that can initiate switchover.

If you are using...	Use this procedure...
Tiebreaker	Removing MetroCluster Configurations in the Tiebreaker Software 1.21 Installation and Configuration Guide
Mediator	Issue the following command from the ONTAP prompt: <code>metrocluster configuration-settings mediator remove</code>
Third-party applications	Refer to the product documentation.

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify Fujitsu support that maintenance is underway. Informing Fujitsu 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.

Procedure

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`

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`

3. Repeat these steps on the partner site.

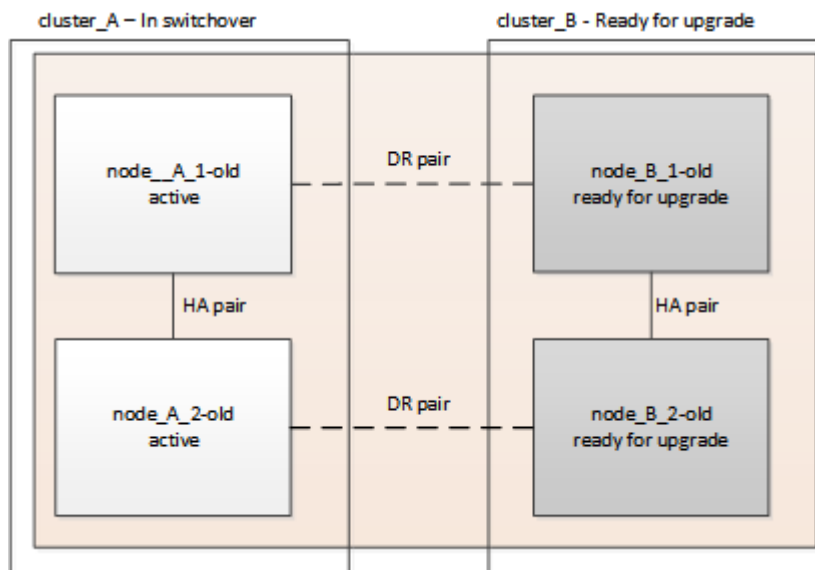
Switching over the MetroCluster configuration

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.



Procedure

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.

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](#) on page 9.

Procedure

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.
 - 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 our 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 elc-80
```
 - b) Remove physical ports from the interface groups: `network port ifgrp remove-port -node node-name -ifgrp interface-group-name -port portid`
For example:

```
network port ifgrp remove-port -node node1 -ifgrp ala -port e0d
```
 - c) Remove VLAN and interface group ports from broadcast domain: `network port broadcast-domain remove-ports -ipSPACE ipSPACE -broadcast-domain broadcast-domain-name -ports nodename:portname, nodename:portname, ..`
 - d) Modify interface group ports to use other physical ports as member as needed: `ifgrp add-port -node node-name -ifgrp interface-group-name -port port-id`
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.

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.

Procedure

1. 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 Guide](#).
 - 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

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.

Setting up the new controllers

You must rack and cable the new controllers.

Procedure

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.
4. Cable the controllers to the IP switches as described in the [MetroCluster IP Installation and Configuration Guide](#).
 - Cabling the IP switches
5. Power up the new nodes and boot them to Maintenance mode.

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.

Procedure

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

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

2. Exit Maintenance mode: `halt`

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

3. Boot the node back into Maintenance mode to enable the configuration changes to take effect:
boot_ontap maint
4. Verify the changes you made:

If you have this type of HBA...	Use this command...
CNA	ucadmin show
FC	fcadmin show

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.

Procedure

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 mcccipha-config modify chassis mcccip`

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](#) on page 13.

Procedure

1. If the nodes being upgraded are AX4100, HX6100 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-idsetenv 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`

Note: 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`

Note: 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`

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.

Procedure

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) Pool11 PZHYNOMD        node_B_1-old(537403322) node_B_1-
old(537403322)
prod4-rk18:9.126L49 node_B_1-old(537403322) Pool11 PPG3J5HA        node_B_1-old(537403322) node_B_1-
old(537403322)
prod4-rk18:8.126L21 node_B_1-old(537403322) Pool11 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.

Using ADP?	Command to use
Yes	<code>disk reassign -s <i>old-sysid</i> -d <i>new-sysid</i> -r <i>dr-partner-sysid</i></code>
No	<code>disk reassign -s <i>old-sysid</i> -d <i>new-sysid</i></code>

4. Reassign the root aggregate disks on the drive shelves to the new controllers: `disk reassign -s old-sysid -d new-sysid`

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

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

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

5. Verify that the disks of the root aggregate are properly reassigned: `disk showstorage`
`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 PZHYNOMD        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      mirror_resync_priority=high(fixed)
                   fast zeroed
                   64-bit
```

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.

Procedure

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, issue the boot menu command for your key management configuration.

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>

6. From the boot menu, select (6) Update flash from backup config.

Note: 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](#) on page 18.
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, again issue the boot menu command for your key management configuration.

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 may need to issue the `recover_XXXXXXXX_keymanager` command and option 6 at the boot menu prompt multiple times until the nodes fully 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> For more information, see Encryption Power Guide .
External key management	<code>security key-manager external restore -vserver SVM -node node -key-server host_name/</code>

If you are using...	Use this command...
	<pre>IP_address:port -key-id key_id -key-tag key_tag node-name</pre>
	<p>For more information, see Encryption Power Guide .</p>

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

a) View the broadcast domains: `network port broadcast-domain show`

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

[Network Management Guide](#)

c) Recreate VLANs and interface groups as needed.

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

[Network Management Guide](#)

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](#) on page 9.

Procedure

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`

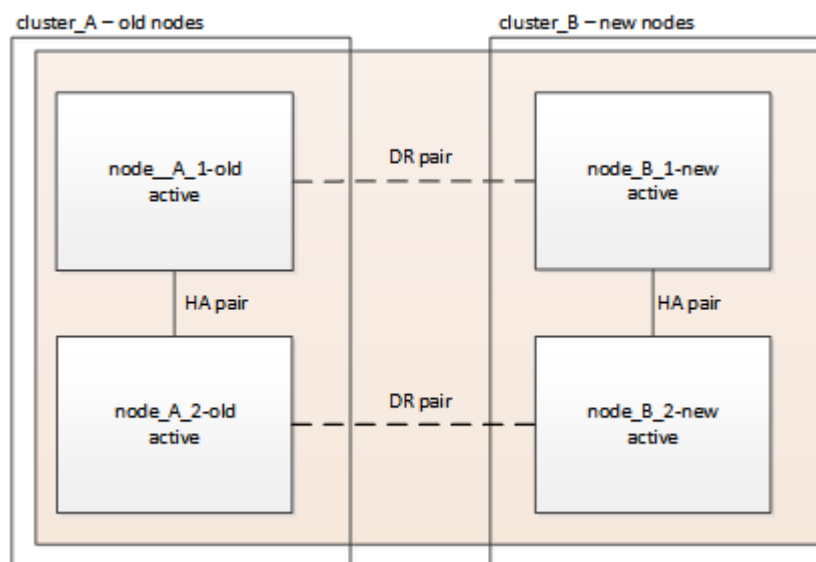
2. Revert the interfaces to their home node: `network interface revert * -vserver vserver_name`

Perform this step on all SVMs as required.

Switching back the MetroCluster configuration

About this task

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.



Procedure

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.

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.

Procedure

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`

Upgrading the nodes on cluster_A

You must repeat the upgrade tasks on cluster_A.

Procedure

Repeat the steps to upgrade the nodes on cluster_A, beginning with [Preparing for the upgrade](#) on page 8.

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.

Restoring Tiebreaker or Mediator monitoring

After completing the upgrade of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

Procedure

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

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations in the MetroCluster Tiebreaker Installation and Configuration Guide
Mediator	Configuring the ONTAP Mediator service from a MetroCluster IP configuration in the MetroCluster IP Installation and Configuration Guide
Third-party applications	Refer to the product documentation.

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.

Procedure

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.

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.

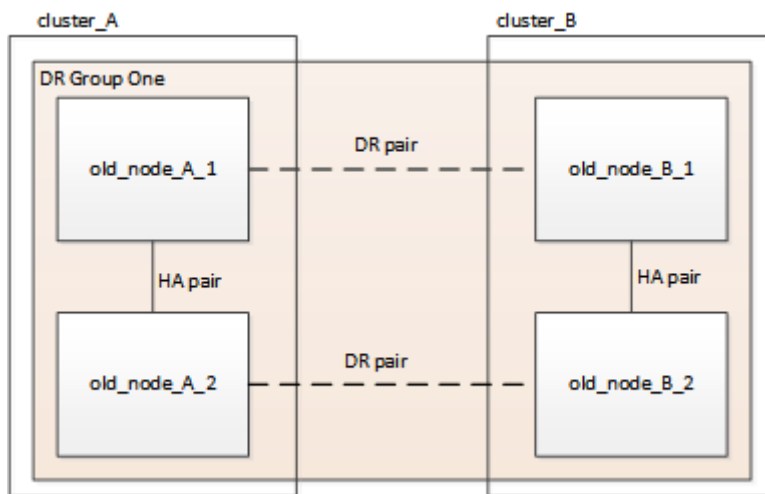
About this task

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

Procedure

1. Gather information from the old nodes.

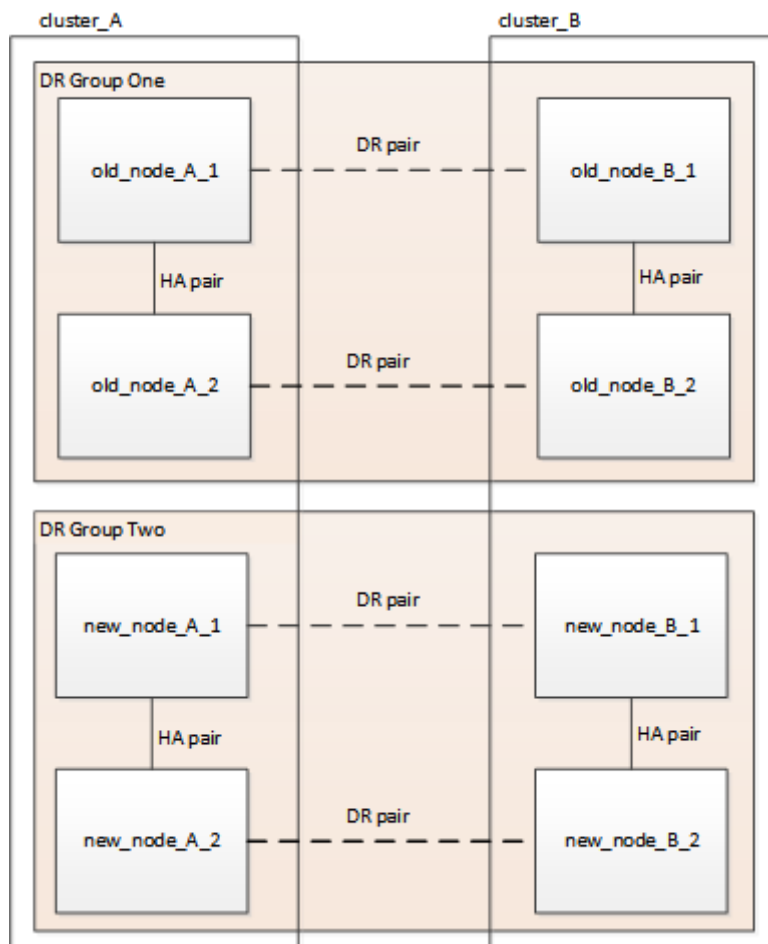
At this stage, the four-node configuration appears as shown in the following image:



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](#) on page 103

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" in [MetroCluster® Service Guide](#).

4. Move the data from the old nodes to new nodes using the following three procedures from the *Controller Hardware Upgrade Express Guide*.

a) Perform all the steps in Creating an aggregate and moving volumes to the new nodes.

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

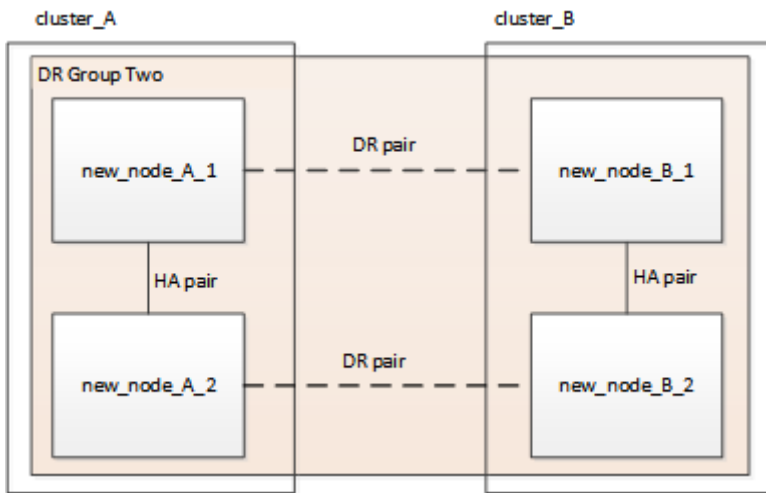
b) Perform all the steps in Moving non-SAN data LIFs and cluster management LIFs to the new nodes.

c) Perform all the steps in Deleting SAN LIFs from the original nodes.

5. Follow the steps in the procedure for removing the old DR group.

[Removing a Disaster Recovery group](#) on page 124

Once you have removed the old DR group (DR group one), the configuration appears as shown in the following image:



Refreshing a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Starting 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 disaster recovery (DR) group.

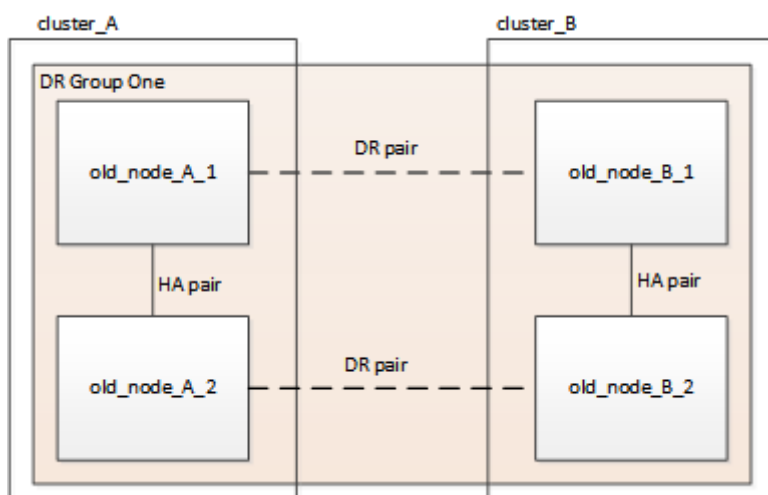
About this task

- This procedure is supported on systems running ONTAP 9.8 and later.
- If you are upgrading the IP switches, they should be upgraded prior to performing this refresh procedure.
- References to "old nodes" mean the nodes that you intend to replace.

Procedure

1. Gather information from the old nodes.

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



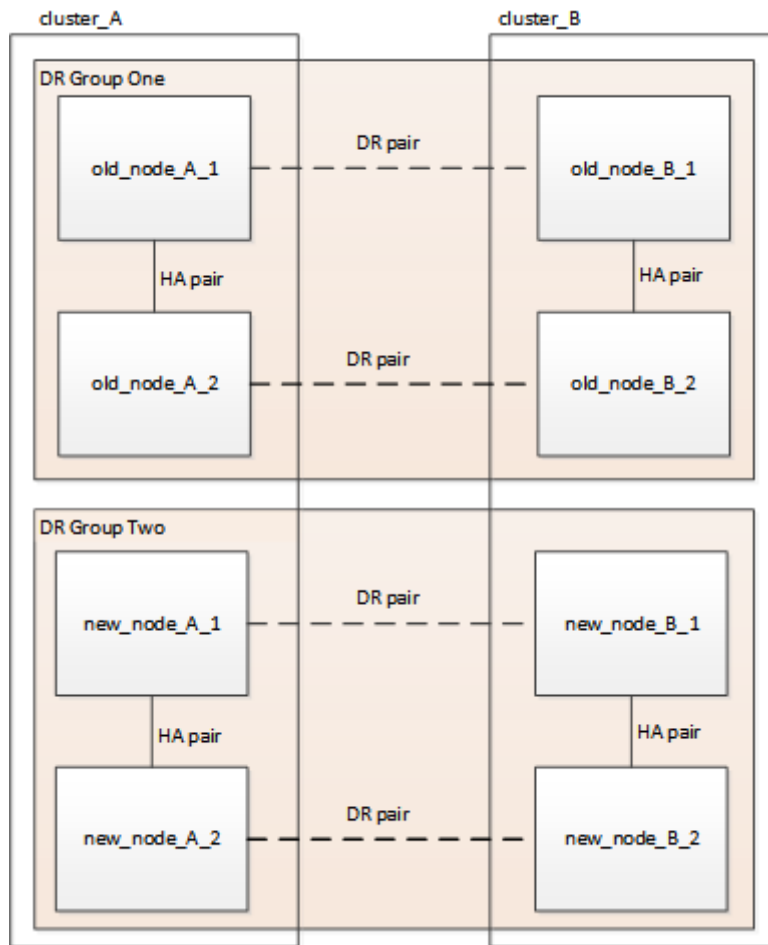
2. To prevent automatic support case generation, send an Autosupport message to indicate the upgrade is underway.
 - a) Issue the following command: `system node autosupport invoke -node * -type all -message "MAINT=10h Upgrading old-model to new-model"`
 Our example specifies a 10 hour maintenance window. You may 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	Removing MetroCluster Configurations in the Tiebreaker Software 1.21 Installation and Configuration Guide
Mediator	Issue the following command from the ONTAP prompt: <code>metrocluster configuration-settings mediator remove</code>
Third-party applications	Refer to the product documentation.

4. Perform all of the steps in the procedure to add nodes to your MetroCluster IP configuration.

[Adding new nodes to the MetroCluster IP configuration](#) on page 31

When the expansion procedure is complete, the configuration appears as shown in the following image:



5. Move the CRS volumes.

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

6. Move the data from the old nodes to new nodes.

- a) Perform all the steps in [Creating an aggregate and moving volumes to the new nodes](#).

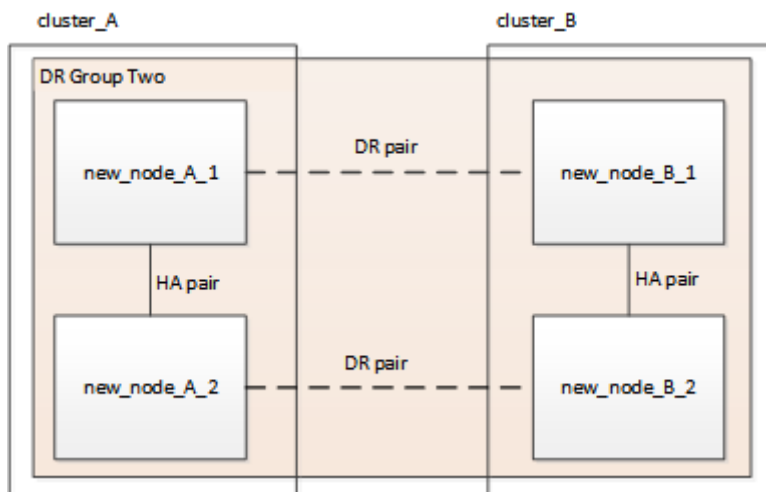
Note: You might choose to mirror the aggregate when or after it is created.

- b) Perform all the steps in [Moving non-SAN data LIFs and cluster management LIFs to the new nodes](#).

7. Follow the steps in the procedure for removing the old DR group.

[Removing a Disaster Recovery group](#) on page 124

Once you have removed the old DR group (DR group one), the configuration appears as shown in the following image:



8. 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`
9. Restore monitoring if necessary, using the procedure for your configuration.

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations in the MetroCluster Tiebreaker Installation and Configuration Guide
Mediator	Configuring the ONTAP Mediator service from a MetroCluster IP configuration in the MetroCluster IP Installation and Configuration Guide
Third-party applications	Refer to the product documentation.

10. 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.

Adding new nodes to the MetroCluster IP configuration

You can add four new nodes to the MetroCluster IP configuration as a second DR group. This creates a temporary eight-node MetroCluster configuration.

Before you begin

- You must complete the initial steps in [Refreshing a four-node MetroCluster IP configuration \(ONTAP 9.8 and later\)](#) on page 29.

After completing this procedure, return to [Refreshing a four-node MetroCluster IP configuration \(ONTAP 9.8 and later\)](#) on page 29 to complete the refresh of the configuration.

- The old and new nodes must be running the same version of ONTAP.
- 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.
- 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.

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"> node_A_1-old node_A_2-old 	<ul style="list-style-type: none"> node_B_1-old node_B_2-old
dr_group_2-new	<ul style="list-style-type: none"> node_A_3-new node_A_4-new 	<ul style="list-style-type: none"> node_B_3-new node_B_4-new

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify Fujitsu support that maintenance is underway. Informing Fujitsu 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.

Procedure

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"`

Our example specifies a 10 hour maintenance window. You may want to allow additional time depending on your plan.

If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period: `system node autosupport invoke -node * -type all -message MAINT=end`

- b) Repeat the command on the partner cluster.

Verifying the health of the MetroCluster configuration

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

Procedure

- Verify the operation of the MetroCluster configuration in ONTAP:
 - Check whether the system is multipathed: `node run -node node-name sysconfig -a`
 - Check for any health alerts on both clusters: `system health alert show`
 - Confirm the MetroCluster configuration and that the operational mode is normal: `metrocluster show`
 - Perform a MetroCluster check: `metrocluster check run`
 - Display the results of the MetroCluster check: `metrocluster check show`
 - Run Config Advisor.
 - After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
- Verify that the cluster is healthy: `cluster show -vserver Cluster`

```
cluster_A::> cluster show -vserver Cluster
Node           Health Eligibility  Epsilon
-----
node_A_1      true   true           false
node_A_2      true   true           false
cluster_A::>
```


3. Verify that all cluster ports are up: `network port show -ipspace cluster`

```
cluster_A::> network port show -ipspace cluster
Node: node_A_1-old

```

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

```
4 entries were displayed.
cluster_A::>
```

4. Verify that all cluster LIFs are up and operational: `network interface show -vserver Cluster`
Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

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

```

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

```
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

```

node_A_2-old_clus2
    true

4 entries were displayed.

cluster_A::>

```

Removing the configuration from monitoring applications

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

Procedure

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

If you are using...	Use this procedure...
Tiebreaker	Removing MetroCluster Configurations in the Tiebreaker Software 1.21 Installation and Configuration Guide
Mediator	Issue the following command from the ONTAP prompt: metrocluster configuration-settings mediator remove
Third-party applications	Refer to the product documentation.

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

Refer to the documentation for the application.

Preparing the new controller modules

You must prepare the four new MetroCluster nodes and install the correct ONTAP version.

About this task

This task must be performed on each of the new nodes:

- node_A_3-new
- node_A_4-new
- node_B_3-new
- node_B_4-new

In these steps, you clear the configuration on the nodes and clear the mailbox region on new drives.

Procedure

1. Rack the new controllers.
2. Cable the new MetroCluster IP nodes to the IP switches as shown in the [MetroCluster Installation and Configuration Guide](#).

Cabling the IP switches

3. Configure the MetroCluster IP nodes using the following sections of the [MetroCluster Installation and Configuration Guide](#).

- a) Gathering required information
- b) Clearing the configuration on a controller module
- c) Verifying the ha-config state of components
- d) Manually assigning drives for pool 0 (ONTAP 9.7 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.

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.

Procedure

1. Add the new MetroCluster IP nodes to the existing MetroCluster configuration.
 - a) Join the first new MetroCluster IP node (node_A_1-new) to the existing MetroCluster IP configuration.

```

Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
  Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to Fujitsu
  Technical
  Support. To disable this feature, enter
  autosupport modify -support disable
  within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
  resolution, should a problem occur on your system.
  For further information on AutoSupport, see:
  http://support.fujitsu.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]: 172.17.8.93

172.17.8.93 is not a valid port.

The physical port that is connected to the node management network.
  Examples of
  node management ports are "e4a" or "e0M".

You can type "back", "exit", or "help" at any question.

Enter the node management interface port [e0M]:
Enter the node management interface IP address: 172.17.8.93
Enter the node management interface netmask: 255.255.254.0
Enter the node management interface default gateway: 172.17.8.1
A node management interface on port e0M with IP address 172.17.8.93 has
  been created.

Use your web browser to complete cluster setup by accessing
  https://172.17.8.93

Otherwise, press Enter to complete cluster setup using the command line
  interface:

Do you want to create a new cluster or join an existing cluster?
  {create, join}:
  join

```

```
Existing cluster interface configuration found:
```

```
Port      MTU      IP              Netmask
e0c       9000     169.254.148.217 255.255.0.0
e0d       9000     169.254.144.238 255.255.0.0
```

```
Do you want to use this configuration? {yes, no} [yes]: yes
```

```
.
.
.
```

b) Join the second new MetroCluster IP node (node_A_2-new) to the existing MetroCluster IP configuration.

2. Repeat these steps to join node_B_1-new and node_B_2-new to cluster_B.

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

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

About this task

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

Procedure

1. On the new MetroCluster IP nodes, configure the intercluster LIFs using the procedures in the [MetroCluster IP Installation and Configuration Guide](#).

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 [MetroCluster IP Installation and Configuration Guide](#).

Considerations for MetroCluster IP configuration

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_A_2-old	node_B_1-old node_B_2-old
	cluster_B	node_B_1-old node_B_2-old	node_A_1-old node_A_2-old
2	cluster_A	node_A_1-new node_A_2-new	node_B_1-new node_B_2-new
	cluster_B	node_B_1-new node_B_2-new	node_A_1-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

Note: 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 ela -address 172.17.26.10 -netmask 255.255.255.0
[Job 260] Job succeeded: Interface Create is successful.

cluster_A::> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_1-new -home-port elb -address 172.17.27.10 -netmask 255.255.255.0
[Job 261] Job succeeded: Interface Create is successful.

cluster_A::> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_2-new -home-port ela -address 172.17.26.11 -netmask 255.255.255.0
[Job 262] Job succeeded: Interface Create is successful.

cluster_A::> :metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_2-new -home-port elb -address 172.17.27.11 -netmask 255.255.255.0
[Job 263] Job succeeded: Interface Create is successful.

cluster_A::> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_1-new -home-port ela -address 172.17.26.12 -netmask 255.255.255.0
[Job 264] Job succeeded: Interface Create is successful.

cluster_A::> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_1-new -home-port elb -address 172.17.27.12 -netmask 255.255.255.0
[Job 265] Job succeeded: Interface Create is successful.

cluster_A::> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_2-new -home-port ela -address 172.17.26.13 -netmask 255.255.255.0
[Job 266] Job succeeded: Interface Create is successful.

cluster_A::> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_2-new -home-port elb -address 172.17.27.13 -netmask 255.255.255.0
[Job 267] Job succeeded: Interface Create is successful.
```

6. Verify the MetroCluster IP interfaces are created: metrocluster configuration-settings interface show

```
cluster_A::>metrocluster configuration-settings interface show
```

DR Group	Cluster	Node	Network Address	Netmask	Gateway	Config State		
1	cluster_A	node_A_1-old	Home Port: ela	172.17.26.10	255.255.255.0	-	completed	
		Home Port: elb	172.17.27.10	255.255.255.0	-	completed		
	node_A_2-old	Home Port: ela	172.17.26.11	255.255.255.0	-	completed		
		Home Port: elb	172.17.27.11	255.255.255.0	-	completed		
	cluster_B	node_B_1-old	Home Port: ela	172.17.26.13	255.255.255.0	-	completed	
		Home Port: elb	172.17.27.13	255.255.255.0	-	completed		
	node_B_1-old	Home Port: ela	172.17.26.12	255.255.255.0	-	completed		
		Home Port: elb	172.17.27.12	255.255.255.0	-	completed		
	2	cluster_A	node_A_3-new	Home Port: ela	172.17.28.10	255.255.255.0	-	completed
			Home Port: elb	172.17.29.10	255.255.255.0	-	completed	
node_A_3-new		Home Port: ela	172.17.28.11	255.255.255.0	-	completed		
		Home Port: elb						

```

cluster_B      172.17.29.11  255.255.255.0  -      completed
  node_B_3-new
    Home Port: ela
      172.17.28.13  255.255.255.0  -      completed
    Home Port: elb
      172.17.29.13  255.255.255.0  -      completed
  node_B_3-new
    Home Port: ela
      172.17.28.12  255.255.255.0  -      completed
    Home Port: elb
      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

Note: 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: ela	172.17.28.10	172.17.28.11	HA Partner	completed	
		Home Port: ela	172.17.28.10	172.17.28.12	DR Partner	completed	
		Home Port: ela	172.17.28.10	172.17.28.13	DR Auxiliary	completed	
		Home Port: elb	172.17.29.10	172.17.29.11	HA Partner	completed	
		Home Port: elb	172.17.29.10	172.17.29.12	DR Partner	completed	
		Home Port: elb	172.17.29.10	172.17.29.13	DR Auxiliary	completed	
		node_A_2-old					
		Home Port: ela	172.17.28.11	172.17.28.10	HA Partner	completed	
		Home Port: ela	172.17.28.11	172.17.28.13	DR Partner	completed	
		Home Port: ela	172.17.28.11	172.17.28.12	DR Auxiliary	completed	
		Home Port: elb	172.17.29.11	172.17.29.10	HA Partner	completed	
		Home Port: elb	172.17.29.11	172.17.29.13	DR Partner	completed	
		Home Port: elb	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: ela	172.17.28.13	172.17.28.12	HA Partner	completed	
		Home Port: ela	172.17.28.13	172.17.28.11	DR Partner	completed	
		Home Port: ela	172.17.28.13	172.17.28.10	DR Auxiliary	completed	
		Home Port: elb	172.17.29.13	172.17.29.12	HA Partner	completed	
		Home Port: elb	172.17.29.13	172.17.29.11	DR Partner	completed	
		Home Port: elb	172.17.29.13	172.17.29.10	DR Auxiliary	completed	
		node_B_1-old					
		Home Port: ela	172.17.28.12	172.17.28.13	HA Partner	completed	
		Home Port: ela	172.17.28.12	172.17.28.10	DR Partner	completed	
		Home Port: ela	172.17.28.12	172.17.28.11	DR Auxiliary	completed	
		Home Port: elb	172.17.29.12	172.17.29.13	HA Partner	completed	
		Home Port: elb	172.17.29.12	172.17.29.10	DR Partner	completed	
		Home Port: elb	172.17.29.12	172.17.29.11	DR Auxiliary	completed	
DR Group	Cluster	Node	Source Network Address	Destination Network Address	Partner	Type	Config State
2	cluster_A						

```

node_A_1-new
  Home Port: ela
    172.17.26.10 172.17.26.11 HA Partner completed
  Home Port: ela
    172.17.26.10 172.17.26.12 DR Partner completed
  Home Port: ela
    172.17.26.10 172.17.26.13 DR Auxiliary completed
  Home Port: elb
    172.17.27.10 172.17.27.11 HA Partner completed
    172.17.27.10 172.17.27.12 DR Partner completed
  Home Port: elb
    172.17.27.10 172.17.27.13 DR Auxiliary completed
node_A_2-new
  Home Port: ela
    172.17.26.11 172.17.26.10 HA Partner completed
    172.17.26.11 172.17.26.13 DR Partner completed
  Home Port: ela
    172.17.26.11 172.17.26.12 DR Auxiliary completed
  Home Port: elb
    172.17.27.11 172.17.27.10 HA Partner completed
    172.17.27.11 172.17.27.13 DR Partner completed
  Home Port: elb
    172.17.27.11 172.17.27.12 DR Auxiliary completed
DR
Group Cluster Node Source Network Address Destination Network Address Partner Type Config State
-----
2 cluster_B
  node_B_2-new
    Home Port: ela
      172.17.26.13 172.17.26.12 HA Partner completed
    Home Port: ela
      172.17.26.13 172.17.26.11 DR Partner completed
    Home Port: ela
      172.17.26.13 172.17.26.10 DR Auxiliary completed
    Home Port: elb
      172.17.27.13 172.17.27.12 HA Partner completed
      172.17.27.13 172.17.27.11 DR Partner completed
    Home Port: elb
      172.17.27.13 172.17.27.10 DR Auxiliary completed
  node_B_1-new
    Home Port: ela
      172.17.26.12 172.17.26.13 HA Partner completed
    Home Port: ela
      172.17.26.12 172.17.26.10 DR Partner completed
    Home Port: ela
      172.17.26.12 172.17.26.11 DR Auxiliary completed
    Home Port: elb
      172.17.27.12 172.17.27.13 HA Partner completed
      172.17.27.12 172.17.27.10 DR Partner completed
    Home Port: elb
      172.17.27.12 172.17.27.11 DR Auxiliary completed
48 entries were displayed.
cluster_A::>

```

9. Verify disk autoassignment and partitioning: `disk show -pool Pool1`

```

cluster_A::> disk show -pool Pool1
Disk Usable Size Shelf Bay Disk Type Container Type Container Name Owner
-----
1.10.4 - 10 4 SAS remote - node_B_2
1.10.13 - 10 13 SAS remote - node_B_2
1.10.14 - 10 14 SAS remote - node_B_1
1.10.15 - 10 15 SAS remote - node_B_1
1.10.16 - 10 16 SAS remote - node_B_1
1.10.18 - 10 18 SAS remote - node_B_2
...
2.20.0 546.9GB 20 0 SAS aggregate aggr0_rhal_a1 node_a_1
2.20.3 546.9GB 20 3 SAS aggregate aggr0_rhal_a2 node_a_2
2.20.5 546.9GB 20 5 SAS aggregate rhal_a1_aggr1 node_a_1
2.20.6 546.9GB 20 6 SAS aggregate rhal_a1_aggr1 node_a_1
2.20.7 546.9GB 20 7 SAS aggregate rhal_a2_aggr1 node_a_2
2.20.10 546.9GB 20 10 SAS aggregate rhal_a1_aggr1 node_a_1
...
43 entries were displayed.
cluster_A::>

```

10. Mirror the root aggregates: `storage aggregate mirror -aggregate aggr0_node_A_1-new`

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

```

Finalizing the addition of the new nodes

You must incorporate the new DR group into the MetroCluster configuration and create mirrored data aggregates on the new nodes.

Procedure

1. 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`

Note: 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)

Physical                               Usable
Size      Position  Disk                Type                Size
-----
-          dparity   5.10.15            SAS                  -
-          parity    5.10.16            SAS                  -
-

```



```

547.1GB data 5.10.17 SAS 546.9GB
558.9GB data 5.10.18 SAS 546.9GB
558.9GB data 5.10.19 SAS 546.9GB

Second Plex

RAID Group rg0, 5 disks (block checksum, raid_dp)
Physical Usable
Size Position Disk Type Size
-----
- dparity 4.20.17 SAS -
- parity 4.20.14 SAS -
547.1GB 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

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::>

```

2. 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`

3. Verify that the nodes are added to their DR group.

```

cluster_A::*> metrocluster node show

DR
Group Cluster Node Configuration State DR Mirroring Mode
-----
1 cluster_A
node_A_1-old configured enabled normal
node_A_2-old configured enabled normal
cluster_B
node_B_1-old configured enabled normal
node_B_2-old configured enabled normal
2 cluster_A
node_A_3-new configured enabled normal
node_A_4-new configured enabled normal
cluster_B
node_B_3-new configured enabled normal
node_B_4-new configured enabled normal
8 entries were displayed.

cluster_A::*>

```

4. Move the MDV_CRS volumes from the old nodes to the new nodes in advanced privilege.

- a) Display the volumes to identify the MDV volumes:

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

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

```
cluster_A::> volume show
Vserver      Volume      Aggregate      State      Type      Size  Available Used%
-----
...
cluster_A    MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_A
              aggr_b1      -              -              RW              -      -      -
cluster_A    MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_B
              aggr_b2      -              -              RW              -      -      -
cluster_A    MDV_CRS_d6b0b313ff5611e9837100a098544e51_A
              aggr_a1      online         RW              10GB           9.50GB  0%
cluster_A    MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
              aggr_a2      online         RW              10GB           9.50GB  0%
...
11 entries were displayed.mple
```

b) Set the advanced privilege level: `set -privilege advanced`

c) Move the MDV volumes, one at a time: `volume move start -volume mdv-volume -destination-aggregate aggr-on-new-node -vserver vservice-name`

The following example shows the command and output for moving

MDV_CRS_d6b0b313ff5611e9837100a098544e51_A to aggregate data_a3 on node_A_3.

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

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

d) Use the `volume show` command to check that the MDV volume has been successfully moved:

```
volume show mdv-name
```

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

```
cluster_A::> vol show MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
Vserver      Volume      Aggregate      State      Type      Size
Available Used%
-----
-----
cluster_A    MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
              aggr_a2      online         RW              10GB
9.50GB      0%
```

e) Return to admin mode: `set -privilege admin`

5. Move epsilon from an old node to a new node:

a) Identify which node currently has epsilon: `cluster show -fields epsilon`

```
cluster_B::> cluster show -fields epsilon
node      epsilon
-----
node_A_1-old    true
node_A_2-old    false
node_A_3-new    false
node_A_4-new    false
4 entries were displayed.
```

b) Set epsilon to false on the old node (node_A_1-old): `cluster modify -node old-node -epsilon false`

c) Set epsilon to true on the new node (node_A_3-new): `cluster modify -node new-node -epsilon true`

d) Verify that epsilon has moved to the correct node: `cluster show -fields epsilon`

```
cluster_A::> cluster show -fields epsilon
```

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

6. Complete the remaining steps in [Refreshing a four-node MetroCluster IP configuration \(ONTAP 9.8 and later\)](#) on page 29.

You have now completed step 4, adding the new nodes to the MetroCluster IP configuration.

Transitioning nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later)

Starting with ONTAP 9.8, nondisruptive transition of workloads and data from an existing four-node MetroCluster FC configuration to a new MetroCluster IP configuration is supported.

- This procedure is supported on systems running ONTAP 9.8 and later.
- This procedure is nondisruptive.

The MetroCluster configuration can continue to serve data during the operation.

- This procedure applies only to a four-node MetroCluster FC configuration.
- You must meet all requirements and follow all steps in the procedure.

Preparing for transition from a MetroCluster FC to a MetroCluster IP configuration

As you prepare for the MetroCluster transition, you should understand the requirements and the steps involved.

Supported platforms for nondisruptive transition

Then transitioning to a MetroCluster IP configuration, you must have a combination of supported platform models.

The following table shows the supported platform combinations. You can transition from platforms in the left-hand column to platforms listed as supported in the columns to the right, as indicated by the green table cells.

Requirements for nondisruptive FC-to-IP transition

Before starting the transition process, you must make sure the configuration meets the requirements.

- It must be a four-node configuration and all nodes must be running ONTAP 9.8 or later.
- The existing and new platforms must be a supported combination for transition.

[Supported platforms for nondisruptive transition](#) on page 44

- It must support a switched cluster configuration.

How transition impacts the MetroCluster hardware components

After completing the transition procedure, key components of the existing MetroCluster configuration have been replaced or reconfigured.

Controller modules

The existing controller modules are replaced by new controller modules. The existing controller modules are decommissioned at the end of the transition procedures.

Storage shelves

Data is moved from the old shelves to the new shelves. The old shelves are decommissioned at the end of the transition procedures.

MetroCluster (back-end) and cluster switches

The back-end switch functionality is replaced by the IP switch fabric. If the MetroCluster FC configuration included FC switches, they are decommissioned at the end of this procedure.

If the MetroCluster FC configuration used cluster switches for the cluster interconnect, in some cases they can be reused to provide the back-end IP switch fabric. Reused

cluster switches must be reconfigured with platform and switch-specific RCFs. procedures.

If the MetroCluster FC configuration did not use cluster switches, new IP switches are added to provide the backend switch fabric.

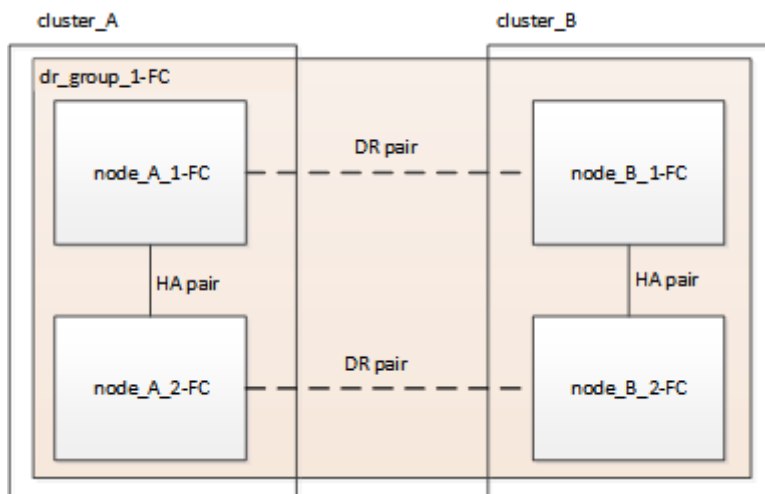
Cluster peering network

The existing customer-provided cluster peering network can be used for the new MetroCluster IP configuration. Cluster peering is configured on the MetroCluster IP nodes as part of the transition procedure.

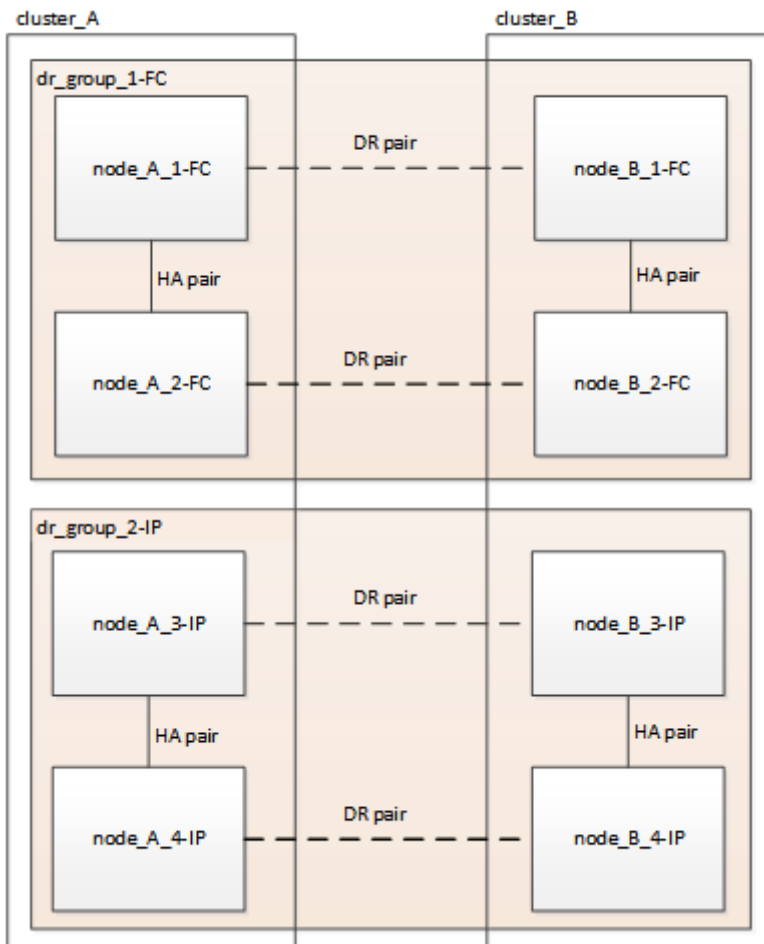
Workflow for nondisruptive MetroCluster transition

You must follow the specific workflow to ensure a successful nondisruptive transition.


The transition process begins with a healthy four-node MetroCluster FC configuration.



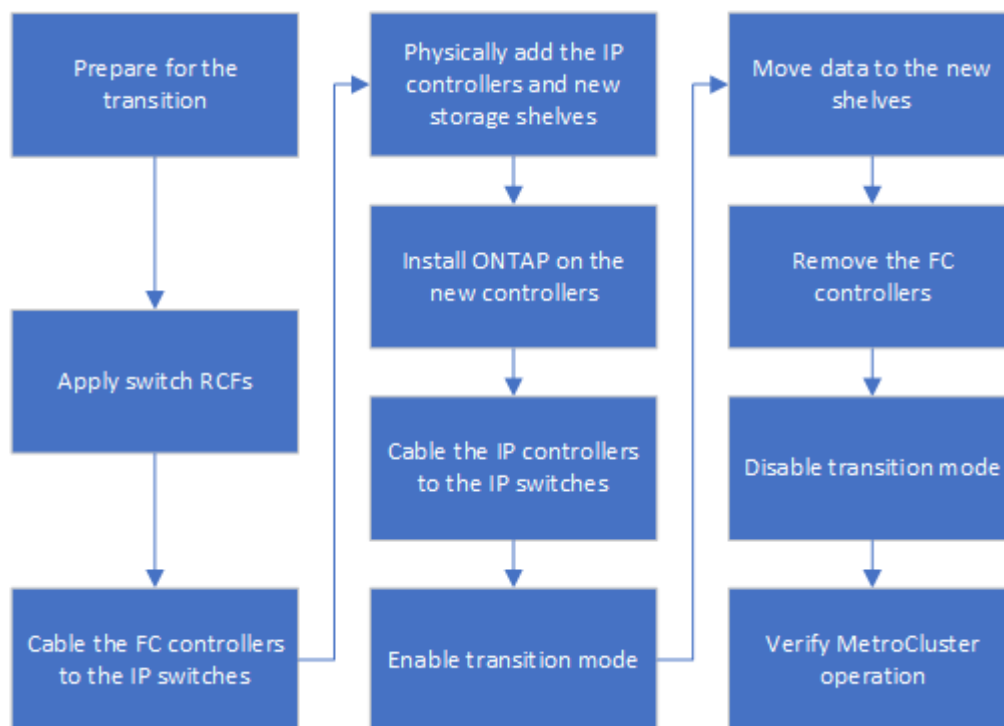
The new MetroCluster IP nodes are added as a second DR group.



Data is transferred from the old DR group to the new DR group, and then the old nodes and their storage are removed from the configuration and decommissioned. The process ends with a four-node MetroCluster IP configuration.



You will use the following workflow to transition the MetroCluster configuration.



Considerations for IP switches

You must ensure the IP switches are supported. If the existing switch model is supported by both the original MetroCluster FC configuration and the new MetroCluster IP configuration, you can reuse the existing switches.

Supported switches

You must use Fujitsu-provided switches.

- The use of MetroCluster-compliant switches (switches that are not validated and provided by Fujitsu) is not supported for transition.
- The IP switches must be supported as a cluster switch by both the MetroCluster FC configuration and the MetroCluster IP configuration.
- The IP switches can be reused in the new MetroCluster IP configuration if the MetroCluster FC is a switched cluster and the IP cluster switches are supported by the MetroCluster IP configuration.
- New IP switches are usually used in the following cases:
 - The MetroCluster FC is a switchless cluster, so new switches are required.
 - The MetroCluster FC is a switched cluster but the existing IP switches are not supported in the MetroCluster IP configuration.
 - You want to use different switches for the MetroCluster IP configuration.

Switchover, healing, and switchback operations during transition

Depending on the stage of the transition process, the MetroCluster switchover, healing, and switchback operations use either the MetroCluster FC or MetroCluster IP workflow.

The following table shows what workflows are used at different stages of the transition process. In some stages, switchover and switchback are not supported.

- In the MetroCluster FC workflow, the switchover, healing and switchback steps are that used by a MetroCluster FC configuration.
- In the MetroCluster IP workflow, the switchover, healing and switchback steps are that used by a MetroCluster IP configuration.
- In the unified workflow, when both the FC and IP nodes are configured, the steps depend on whether NSO or USO is performed. The details are shown in the table.

Note: Automatic unplanned switchover is not available during the transition process.

Stage of transition	Workflow used	
	Negotiated switchover uses...	Unplanned switchover uses...
Before the MetroCluster IP nodes have joined the cluster.	MetroCluster FC	MetroCluster FC
After the MetroCluster IP nodes have joined the cluster, before the <code>metrocluster configure</code> command is performed.	Not supported	MetroCluster FC
After the <code>metrocluster configure</code> command has been issued. Volume move can be in progress.	Unified: All remote site nodes remain up and healing is done automatically	Unified: <ul style="list-style-type: none"> • Mirrored aggregates owned by the MetroCluster FC node are mirrored if storage is accessible, all others are degraded after switchover • All remote site nodes are able to boot up. • Heal aggregate and heal root commands must be run manually.
The MetroCluster FC nodes have been unconfigured.	Not supported	MetroCluster IP
The <code>cluster unjoin</code> command has been performed on the MetroCluster FC nodes.	MetroCluster IP	MetroCluster IP

Alert messages and tool support during transition

You may notice alert messages during transition. These alerts can be safely ignored. Also, some tools are not available during transition.

- ARS may alert during transition.
These alerts can be ignored and should disappear once the transition has finished.
- Active IQ Unified Manager may alert during transition.
These alerts can be ignored and should disappear once the transition has finished.
- Config Advisor is not supported during transition.
- System Manager is not supported during transition.

Example naming in this procedure

This procedure uses example names throughout to identify the DR groups, nodes, and switches involved.

DR groups	cluster_A at site_A	cluster_B at site_B
dr_group_1-FC	<ul style="list-style-type: none"> • node_A_1-FC • node_A_2-FC 	<ul style="list-style-type: none"> • node_B_1-FC • node_B_2-FC
dr_group_2-IP	<ul style="list-style-type: none"> • node_A_3-IP • node_A_4-IP 	<ul style="list-style-type: none"> • node_B_3-IP • node_B_4-IP

DR groups	cluster_A at site_A	cluster_B at site_B
Switches	Initial switches (if fabric-attached configuration): <ul style="list-style-type: none"> switch_A_1-FC switch_A_2-FC MetroCluster IP switches: <ul style="list-style-type: none"> switch_A_1-IP switch_A_2-IP 	Initial switches (if fabric-attached configuration): <ul style="list-style-type: none"> switch_B_1-FC switch_B_2-FC MetroCluster IP switches: <ul style="list-style-type: none"> switch_B_1-IP switch_B_2-IP

Transitioning from MetroCluster FC to MetroCluster IP configurations

After reviewing all requirements and preparing for the transition you perform the transition procedure. This includes physically installing and cabling the MetroCluster IP components, configuring the switches, moving data to the new configuration, and removing the MetroCluster FC controller modules and storage shelves.

Verifying the health of the MetroCluster configuration

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

Procedure

- Verify the operation of the MetroCluster configuration in ONTAP:
 - Check whether the system is multipathed: `node run -node node-name sysconfig -a`
 - Check for any health alerts on both clusters: `system health alert show`
 - Confirm the MetroCluster configuration and that the operational mode is normal: `metrocluster show`
 - Perform a MetroCluster check: `metrocluster check run`
 - Display the results of the MetroCluster check: `metrocluster check show`
 - Check for any health alerts on the switches (if present): `storage switch show`
 - Run Config Advisor.
 - After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.
- Verify that the cluster is healthy: `cluster show`

```
cluster_A::> cluster show
Node           Health  Eligibility  Epsilon
-----
node_A_1_FC    true   true         false
node_A_2_FC    true   true         false
cluster_A::>
```

- Verify that all cluster ports are up: `network port show -ipSPACE cluster`

```
cluster_A::> network port show -ipSPACE cluster
Node: node_A_1_FC

Port           IPspace      Broadcast Domain  Link  MTU      Speed(Mbps) Health
Admin/Oper    Status
-----
e0a            Cluster      Cluster           up    9000    auto/10000 healthy
e0b            Cluster      Cluster           up    9000    auto/10000 healthy
Node: node_A_2_FC

Speed(Mbps) Health
```

Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper	Status
e0a	Cluster	Cluster	up	9000	auto/10000	healthy
e0b	Cluster	Cluster	up	9000	auto/10000	healthy

4 entries were displayed.

cluster_A::>

4. Verify that all cluster LIFs are up and operational: `network interface show -vserver cluster`
Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up

cluster_A::> network interface show -vserver cluster

Current Is Vserver Home	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Port
Cluster	node_A-1_FC_clus1	up/up	169.254.209.69/16	node_A-1_FC	e0a
true	node_A_1_FC_clus2	up/up	169.254.49.125/16	node_A_1_FC	e0b
true	node_A_2_FC_clus1	up/up	169.254.47.194/16	node_A_2_FC	e0a
true	node_A_2_FC_clus2	up/up	169.254.19.183/16	node_A_2_FC	e0b

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_FC_clus1	true
	node_A_1_FC_clus2	true
	node_A_2_FC_clus1	true
	node_A_2_FC_clus2	true

4 entries were displayed.

cluster_A::>

Removing the existing configuration from the Tiebreaker or other monitoring software

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

Procedure

1. Remove the existing MetroCluster configuration from the Tiebreaker software.
Removing MetroCluster configurations
2. Remove the existing MetroCluster configuration from any third-party application that can initiate switchover.
Refer to the documentation for the application.

Generating and applying RCFs to the new IP switches

If you are using new IP switches for the MetroCluster IP configuration, you must configure the switches with a custom RCF file.

About this task

This task is required if you are using new switches.

If you are using existing switches, proceed to [Moving the local cluster connections](#) on page 51.

Procedure

1. Install and rack the new IP switches.
2. Prepare the IP switches for the application of the new RCF files.
Follow the steps in the section for your switch vendor from the [MetroCluster IP Installation and Configuration guide](#).
 - Resetting the Broadcom IP switch to factory defaults
 - Resetting the Cisco IP switch to factory defaults
3. Update the firmware on the switch to a supported version, if necessary.
4. Use the RCF generator tool to create the RCF file depending on your switch vendor and the platform models, and then update the switches with the file.
Follow the steps in the section for your switch vendor from the [MetroCluster IP Installation and Configuration guide](#).
 - Downloading and installing the Broadcom IP RCF files
 - Downloading and installing the Cisco IP RCF files

Moving the local cluster connections

You must move the MetroCluster FC configuration's cluster interfaces to the IP switches.

Moving the cluster connections on the MetroCluster FC nodes

You must move the cluster connections on the MetroCluster FC nodes to the IP switches. The steps depend on whether you are using the existing IP switches or you are using new IP switches.

About this task

You must perform this task on both MetroCluster sites.

The following task assumes a controller module using two ports for the cluster connections. Some controller module models use four or more ports for the cluster connection. In that case, for the purposes of this example, the ports are divided into two groups, alternating ports between the two groups

The following table shows the example ports used in this task.

Number of cluster connections on the controller module	Group A ports	Group B ports
Two	e0a	e0b
Four	e0a, e0c	e0b, e0d

- *group A* ports connect to local switch switch_x_1-IP
- *group B* ports connect to local switch switch_x_2-IP

The following table shows which switch ports the FC nodes connect to. For the Broadcom BES-53248 switch, the port usage depends on the model of the MetroCluster IP nodes.

Switch model	MetroCluster IP node model	Switch port(s)	Connects to
Cisco 3132Q-V or 3232C	Any	5	node_x_1-FC
		6	node_x_2-FC
Broadcom BES-53248	HX2200/AX2100	1, 2, 3	node_x_1-FC
	HX2200/AX2100	4, 5, 6	node_x_1-FC

Moving the local cluster connections when using new IP switches

If you are using new IP switches, you must physically move the existing MetroCluster FC nodes' cluster connections to the new switches.

Procedure

1. Move the MetroCluster FC node group A cluster connections to the new IP switches.
Use the ports described in [Moving the cluster connections on the MetroCluster FC nodes](#) on page 51.
 - a) Disconnect all the group A ports from the switch, or, if the MetroCluster FC configuration was a switchless cluster, disconnect them from the partner node.
 - b) Disconnect the group A ports from node_A_1-FC and node_A_2-FC.
 - c) Connect the group A ports of node_A_1-FC to the switch ports for the FC node on switch_A_1-IP
 - d) Connect the group A ports of node_A_2-FC to the switch ports for the FC node on switch_A_1-IP
2. Verify that all cluster ports are up: `network port show -ipspace Cluster`

```
cluster_A::*> network port show -ipspace Cluster
Node: node_A_1-FC
Port          IPspace      Broadcast Domain  Link  MTU  Speed(Mbps)  Health
-----
Admin/Oper   Status
e0a          Cluster      Cluster           up    9000 auto/10000    healthy
e0b          Cluster      Cluster           up    9000 auto/10000    healthy
Node: node_A_2-FC
Port          IPspace      Broadcast Domain  Link  MTU  Speed(Mbps)  Health
-----
Admin/Oper   Status
e0a          Cluster      Cluster           up    9000 auto/10000    healthy
e0b          Cluster      Cluster           up    9000 auto/10000    healthy
4 entries were displayed.
cluster_A::*>
```

3. Verify that all interfaces display true in the Is Home column: `network interface show -vserver cluster`

This might take several minutes to complete.

```
cluster_A::*> network interface show -vserver cluster
Current Is Logical      Status      Network      Current
Vserver Interface Admin/Oper Address/Mask Node          Port
Home
```

```

-----
Cluster
node_A_1_FC_clus1
up/up      169.254.209.69/16  node_A_1_FC  e0a
true
node_A_1-FC_clus2
up/up      169.254.49.125/16  node_A_1-FC  e0b
true
node_A_2-FC_clus1
up/up      169.254.47.194/16  node_A_2-FC  e0a
true
node_A_2-FC_clus2
up/up      169.254.19.183/16  node_A_2-FC  e0b
true

4 entries were displayed.

cluster_A::*>

```

4. Perform the above steps on both nodes (node_A_1-FC and node_A_2-FC) to move the group B ports of the cluster interfaces.
5. Repeat the above steps on the partner cluster, cluster_B.

Moving the local cluster connections when reusing existing IP switches

If you are reusing existing IP switches, you must update firmware, reconfigure the switches with the correct Reference Configure Files (RCFs) and move the connections to the correct ports one switch at a time.

Before you begin

This task is required only if the FC nodes are connected to existing IP switches and you are reusing the switches.

Procedure

1. Disconnect the local cluster connections that connect to switch_A_1_IP
 - a) Disconnect the group A ports from the existing IP switch.
 - b) Disconnect the ISL ports on switch_A_1_IP.

You can see the Installation and Setup instructions for the platform to see the cluster port usage.

[AX2100/HX2000 Systems Installation and Setup Instructions](#)

[HX6100 Systems Installation and Setup Instructions](#)

2. Reconfigure switch_A_1_IP using RCF files generated for your platform combination and transition. Follow the steps in the section for your switch vendor from the [MetroCluster IP Installation and Configuration guide](#), as given in the links below.
 - a) If required, download and install the new switch firmware.

You should use the latest firmware that the MetroCluster IP nodes support.

 - Downloading and installing the Broadcom switch EFOS software
 - Downloading and installing the Cisco switch NX-OS software
 - b) Prepare the IP switches for the application of the new RCF files.

- Resetting the Broadcom IP switch to factory defaults
 - Resetting the Cisco IP switch to factory defaults
- c) Download and install the IP RCF file depending on your switch vendor.
- Downloading and installing the Broadcom IP RCF files
 - Downloading and installing the Cisco IP RCF files

3. Reconnect the group A ports to switch_A_1_IP.

Use the ports described in [Moving the cluster connections on the MetroCluster FC nodes](#) on page 51.

4. Verify that all cluster ports are up: `network port show -ipSPACE cluster`

```
Cluster-A::*> network port show -ipSPACE cluster

Node: node_A_1_FC

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

Node: node_A_2_FC

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

4 entries were displayed.

Cluster-A::*>
```

5. Verify that all interfaces are on their home port: `network interface show -vserver Cluster`

```
Cluster-A::*> network interface show -vserver Cluster

Vserver      Logical      Status      Network      Current      Current Is
Interface    Admin/Oper   Address/Mask Node          Port         Home
-----
Cluster
node_A_1_FC_clus1 up/up        169.254.209.69/16 node_A_1_FC  e0a         true
node_A_1_FC_clus2 up/up        169.254.49.125/16 node_A_1_FC  e0b         true
node_A_2_FC_clus1 up/up        169.254.47.194/16 node_A_2_FC  e0a         true
node_A_2_FC_clus2 up/up        169.254.19.183/16 node_A_2_FC  e0b         true

4 entries were displayed.

Cluster-A::*>
```

6. Repeat all the previous steps on switch_A_2_IP.
7. Reconnect the local cluster ISL ports.
8. Repeat the above steps at site_B for switch B_1_IP and switch B_2_IP.
9. Connect the remote ISLs between the sites.

Verifying that the cluster connections are moved and the cluster is healthy

To ensure that there is proper connectivity and that the configuration is ready to proceed with the transition process, you must verify that the cluster connections are moved correctly, the cluster switches are recognized and the cluster is healthy.

Procedure

1. Verify that all cluster ports are up and running: `network port show -ipSPACE Cluster`

```
Cluster-A::*> network port show -ipSPACE Cluster

Node: Node-A-1-FC

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

Node: Node-A-2-FC

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

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

4 entries were displayed.

Cluster-A::*>

2. Verify that all interfaces are on their home port: `network interface show -vserver Cluster`
This may take several minutes to complete.

The following example shows that all interfaces show `true` in the `Is Home` column.

Cluster-A::*> network interface show -vserver Cluster

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
Cluster	Node-A-1-FC_clus1	up/up	169.254.209.69/16	Node-A-1-FC	e0a	true
	Node-A-1-FC_clus2	up/up	169.254.49.125/16	Node-A-1-FC	e0b	true
	Node-A-2-FC_clus1	up/up	169.254.47.194/16	Node-A-2-FC	e0a	true
	Node-A-2-FC_clus2	up/up	169.254.19.183/16	Node-A-2-FC	e0b	true

4 entries were displayed.

Cluster-A::*>

3. Verify that both the local IP switches are discovered by the nodes: `network device-discovery show -protocol cdp`

Cluster-A::*> network device-discovery show -protocol cdp

Node/Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform
Node-A-1-FC	/cdp			
	e0a	Switch-A-3-IP	1/5/1	N3K-C3232C
Node-A-2-FC	e0b	Switch-A-4-IP	0/5/1	N3K-C3232C
	/cdp			
Node-A-1-FC	e0a	Switch-A-3-IP	1/6/1	N3K-C3232C
	e0b	Switch-A-4-IP	0/6/1	N3K-C3232C

4 entries were displayed.

Cluster-A::*>

4. On the IP switch, verify that the MetroCluster IP nodes have been discovered by both local IP switches: `show cdp neighbors`

You must perform this step on each switch.

This example shows how to verify the nodes are discovered on Switch-A-3-IP.

(Switch-A-3-IP)# show cdp neighbors

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID         Local Intrfce  Hldtme Capability  Platform      Port ID
Node-A-1-FC       Eth1/5/1      133    H             FAS8200       e0a
Node-A-2-FC       Eth1/6/1      133    H             FAS8200       e0a
Switch-A-4-IP(FD0220329A4)
                  Eth1/7        175    R S I s       N3K-C3232C    Eth1/7
Switch-A-4-IP(FD0220329A4)
                  Eth1/8        175    R S I s       N3K-C3232C    Eth1/8
Switch-B-3-IP(FD0220329B3)
                  Eth1/20       173    R S I s       N3K-C3232C    Eth1/20
Switch-B-3-IP(FD0220329B3)
                  Eth1/21       173    R S I s       N3K-C3232C    Eth1/21
```

Total entries displayed: 4

(Switch-A-3-IP)#

This example shows how to verify the nodes are discovered on Switch-A-4-IP.

(Switch-A-4-IP)# show cdp neighbors

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID         Local Intrfce  Hldtme Capability  Platform      Port ID
Node-A-1-FC       Eth1/5/1      133    H             FAS8200       e0b
Node-A-2-FC       Eth1/6/1      133    H             FAS8200       e0b
```

```

Switch-A-3-IP(FDO220329A3)
    Eth1/7          175  R S I s  N3K-C3232C  Eth1/7
Switch-A-3-IP(FDO220329A3)
    Eth1/8          175  R S I s  N3K-C3232C  Eth1/8
Switch-B-4-IP(FDO220329B4)
    Eth1/20         169  R S I s  N3K-C3232C  Eth1/20
Switch-B-4-IP(FDO220329B4)
    Eth1/21         169  R S I s  N3K-C3232C  Eth1/21

Total entries displayed: 4

(Switch-A-4-IP)#

```

Preparing the MetroCluster IP controllers

You must prepare the four new MetroCluster IP nodes and install the correct ONTAP version.

About this task

This task must be performed on each of the new nodes:

- node_A_1-IP
- node_A_2-IP
- node_B_1-IP
- node_B_2-IP

In these steps, you clear the configuration on the nodes and clear the mailbox region on new drives.

Procedure

1. Rack the new controllers for the MetroCluster IP configuration.
The MetroCluster FC nodes (node_A_x-FC and node_B_x-FC) remain cabled at this time.
2. Cable the MetroCluster IP nodes to the IP switches as shown in the [MetroCluster Installation and Configuration Guide](#).
Cabling the IP switches
3. Configure the MetroCluster IP nodes using the following sections of the *MetroCluster Installation and Configuration Guide*.
 - a) Gathering required information
 - b) Restoring system defaults on a previously used controller module
 - c) Verifying the ha-config state of components
 - d) Manually assigning drives for pool 0 (ONTAP 9.7 and later)
4. From Maintenance mode, issue the `halt` command to exit Maintenance mode, and then issue the `boot_ontap` command to boot the system and get to cluster setup.
Do not complete the cluster wizard or node wizard at this time.
5. Repeat these steps on the other MetroCluster IP nodes.

Configure the MetroCluster for transition

To prepare the configuration for transition you add the new nodes to the existing MetroCluster configuration and then move data to the new nodes.

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify Fujitsu support that maintenance is underway. Informing Fujitsu 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.

Procedure

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.

Enabling transition mode and disabling cluster HA

You must enable the MetroCluster transition mode to allow the old and new nodes to operate together in the MetroCluster configuration, and disable cluster HA.

Procedure

1. Enable transition:

- a) Change to the advanced privilege level: `set -privilege advanced`
- b) Enable transition mode: `metrocluster transition enable -transition-mode non-disruptive`

Note: Run this command on one cluster only.

```
cluster_A::*> metrocluster transition enable -transition-mode non-disruptive

Warning: This command enables the start of a "non-disruptive" MetroCluster
FC-to-IP transition. It allows the addition of hardware for another DR
group that uses IP fabrics, and the removal of a DR group that uses FC
fabrics. Clients will continue to access their data during a
non-disruptive transition.

Automatic unplanned switchover will also be disabled by this command.
Do you want to continue? {y|n}: y

cluster_A::*>
```

- c) Return to the admin privilege level: `set -privilege admin`

2. Verify that transition is enabled on both the clusters.

```
cluster_A::*> metrocluster transition show-mode
Transition Mode
-----
non-disruptive

cluster_A::*>

cluster_B::*> metrocluster transition show-mode
Transition Mode
-----
non-disruptive

Cluster_B::>
```

3. Disable cluster HA.

Note: You must run this command on both clusters.

```
cluster_A::*> cluster ha modify -configured false

Warning: This operation will unconfigure cluster HA. Cluster HA must be
configured on a two-node cluster to ensure data access availability in
the event of storage failover.
Do you want to continue? {y|n}: y
Notice: HA is disabled.

cluster_A::*>

cluster_B::*> cluster ha modify -configured false

Warning: This operation will unconfigure cluster HA. Cluster HA must be
configured on a two-node cluster to ensure data access availability in
the event of storage failover.
Do you want to continue? {y|n}: y
Notice: HA is disabled.

cluster_B::*>
```

4. Verify that cluster HA is disabled.

Note: You must run this command on both clusters.

```
cluster_A::> cluster ha show
```

```

High Availability Configured: false
Warning: Cluster HA has not been configured. Cluster HA must be configured
on a two-node cluster to ensure data access availability in the
event of storage failover. Use the "cluster ha modify -configured
true" command to configure cluster HA.

cluster_A::>

cluster_B::> cluster ha show

High Availability Configured: false
Warning: Cluster HA has not been configured. Cluster HA must be configured
on a two-node cluster to ensure data access availability in the
event of storage failover. Use the "cluster ha modify -configured
true" command to configure cluster HA.

cluster_B::>

```

Joining the MetroCluster IP nodes to the clusters

You must add the four new MetroCluster IP nodes to the existing MetroCluster configuration.

About this task

You must perform this task on both clusters.

Procedure

1. Add the MetroCluster IP nodes to the existing MetroCluster configuration.
 - a) Join the first MetroCluster IP node (node_A_1-IP) to the existing MetroCluster FC configuration.

```

Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
  Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.

This system will send event messages and periodic reports to Fujitsu
  Technical
  Support. To disable this feature, enter autosupport modify -support
  disable
  within 24 hours.

Enabling AutoSupport can significantly speed problem determination and
  resolution, should a problem occur on your system.
  For further information on AutoSupport, see:
  http://support.fujitsu.com/autosupport/

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]:
Enter the node management interface IP address: 172.17.8.93
Enter the node management interface netmask: 255.255.254.0
Enter the node management interface default gateway: 172.17.8.1
A node management interface on port e0M with IP address 172.17.8.93 has
  been created.

Use your web browser to complete cluster setup by accessing
  https://172.17.8.93

Otherwise, press Enter to complete cluster setup using the command line
  interface:

Do you want to create a new cluster or join an existing cluster?
  {create, join}:
  join

```

```
Existing cluster interface configuration found:
```

```
Port      MTU      IP                Netmask
e0c       9000    169.254.148.217  255.255.0.0
e0d       9000    169.254.144.238  255.255.0.0
```

```
Do you want to use this configuration? {yes, no} [yes]: yes
```

```
.
.
.
```

b) Join the second MetroCluster IP node (node_A_2-IP) to the existing MetroCluster FC configuration.

2. Repeat these steps to join node_B_1-IP and node_B_2-IP to cluster_B.

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

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

About this task

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

Procedure

1. On the new MetroCluster IP nodes, configure the intercluster LIFs using the procedures in the [MetroCluster IP Installation and Configuration Guide](#).

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
```

3. Configure the DR group for the MetroCluster IP nodes: `metrocluster configuration-settings dr-group create -partner-cluster`

```
cluster_A::> metrocluster configuration-settings dr-group create -partner-cluster
cluster_B -local-node node_A_3-IP -remote-node node_B_3-IP
[Job 259] Job succeeded: DR Group Create is successful.
cluster_A::>
```

4. Verify that the DR group is created. `metrocluster configuration-settings dr-group show`

```
cluster_A::> metrocluster configuration-settings dr-group show

DR Group ID Cluster      Node                DR Partner Node
-----
2          cluster_A      node_A_3-IP        node_B_3-IP
           cluster_A      node_A_4-IP        node_B_4-IP
           cluster_B      node_B_3-IP        node_A_3-IP
           cluster_B      node_B_4-IP        node_A_4-IP

4 entries were displayed.
cluster_A::>
```

You will notice that the DR group for the old MetroCluster FC nodes (DR Group 1) is not listed when you run the `metrocluster configuration-settings dr-group show` command.

You can use `metrocluster node show` command on both sites to list all nodes.

```
cluster_A:> metrocluster node show
```

DR Group	Cluster	Node	Configuration State	DR Mirroring	Mode	
1	cluster_A	node_A_1-FC	configured	enabled	normal	
		node_A_2-FC	configured	enabled	normal	
	cluster_B	node_B_1-FC	configured	enabled	normal	
		node_B_2-FC	configured	enabled	normal	
	2	cluster_A	node_A_1-IP	ready to configure	-	-
			node_A_2-IP	ready to configure	-	-

```
cluster_B:> metrocluster node show
```

DR Group	Cluster	Node	Configuration State	DR Mirroring	Mode
1	cluster_B	node_B_1-FC	configured	enabled	normal
		node_B_2-FC	configured	enabled	normal
	cluster_A	node_A_1-FC	configured	enabled	normal
		node_A_2-FC	configured	enabled	normal
2	cluster_B	node_B_1-IP	ready to configure	-	-
		node_B_2-IP	ready to configure	-	-

5. Configure the MetroCluster IP interfaces for the newly joined MetroCluster IP nodes: `metrocluster configuration-settings interface create -cluster-name`

Note: You can configure the MetroCluster IP interfaces from either cluster.

```
cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_3-IP -home-port elb -address 172.17.26.10 -netmask 255.255.255.0
[Job 260] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_3-IP -home-port elb -address 172.17.27.10 -netmask 255.255.255.0
[Job 261] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_4-IP -home-port elb -address 172.17.26.11 -netmask 255.255.255.0
[Job 262] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_A -home-node
node_A_4-IP -home-port elb -address 172.17.27.11 -netmask 255.255.255.0
[Job 263] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_3-IP -home-port elb -address 172.17.26.12 -netmask 255.255.255.0
[Job 264] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_3-IP -home-port elb -address 172.17.27.12 -netmask 255.255.255.0
[Job 265] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_4-IP -home-port elb -address 172.17.26.13 -netmask 255.255.255.0
[Job 266] Job succeeded: Interface Create is successful.

cluster_A:> metrocluster configuration-settings interface create -cluster-name cluster_B -home-node
node_B_4-IP -home-port elb -address 172.17.27.13 -netmask 255.255.255.0
[Job 267] Job succeeded: Interface Create is successful.
```

6. Verify the MetroCluster IP interfaces are created: `metrocluster configuration-settings interface show`

```
cluster_A:> metrocluster configuration-settings interface show
```

DR Group	Cluster	Node	Network Address	Netmask	Gateway	Config State	
2	cluster_A	node_A_3-IP	Home Port: elb	172.17.26.10	255.255.255.0	-	completed
			Home Port: elb	172.17.27.10	255.255.255.0	-	completed
		node_A_4-IP	Home Port: elb	172.17.26.11	255.255.255.0	-	completed
			Home Port: elb	172.17.27.11	255.255.255.0	-	completed

```

      Home Port: elb
      172.17.27.11      255.255.255.0 -      completed
cluster_B
  node_B_3-IP
    Home Port: ela
    172.17.26.13      255.255.255.0 -      completed
    Home Port: elb
    172.17.27.13      255.255.255.0 -      completed
  node_B_3-IP
    Home Port: ela
    172.17.26.12      255.255.255.0 -      completed
    Home Port: elb
    172.17.27.12      255.255.255.0 -      completed
8 entries were displayed.
cluster_A>

```

7. Connect the MetroCluster IP interfaces: metrocluster configuration-settings connection connect

Note: 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      Destination
Network Address Network Address Partner Type Config State
-----
2      cluster_A
  node_A_3-IP
    Home Port: ela
    172.17.26.10      172.17.26.11      HA Partner  completed
    Home Port: ela
    172.17.26.10      172.17.26.12      DR Partner  completed
    Home Port: ela
    172.17.26.10      172.17.26.13      DR Auxiliary completed
    Home Port: elb
    172.17.27.10      172.17.27.11      HA Partner  completed
    Home Port: elb
    172.17.27.10      172.17.27.12      DR Partner  completed
    Home Port: elb
    172.17.27.10      172.17.27.13      DR Auxiliary completed
  node_A_4-IP
    Home Port: ela
    172.17.26.11      172.17.26.10      HA Partner  completed
    Home Port: ela
    172.17.26.11      172.17.26.13      DR Partner  completed
    Home Port: ela
    172.17.26.11      172.17.26.12      DR Auxiliary completed
    Home Port: elb
    172.17.27.11      172.17.27.10      HA Partner  completed
    Home Port: elb
    172.17.27.11      172.17.27.13      DR Partner  completed
    Home Port: elb
    172.17.27.11      172.17.27.12      DR Auxiliary completed

DR
Group Cluster Node      Source      Destination
Network Address Network Address Partner Type Config State
-----
2      cluster_B
  node_B_4-IP
    Home Port: ela
    172.17.26.13      172.17.26.12      HA Partner  completed
    Home Port: ela
    172.17.26.13      172.17.26.11      DR Partner  completed
    Home Port: ela
    172.17.26.13      172.17.26.10      DR Auxiliary completed
    Home Port: elb
    172.17.27.13      172.17.27.12      HA Partner  completed
    Home Port: elb
    172.17.27.13      172.17.27.11      DR Partner  completed
    Home Port: elb
    172.17.27.13      172.17.27.10      DR Auxiliary completed
  node_B_3-IP
    Home Port: ela
    172.17.26.12      172.17.26.13      HA Partner  completed
    Home Port: ela
    172.17.26.12      172.17.26.10      DR Partner  completed
    Home Port: ela
    172.17.26.12      172.17.26.11      DR Auxiliary completed
    Home Port: elb
    172.17.27.12      172.17.27.13      HA Partner  completed
    Home Port: elb
    172.17.27.12      172.17.27.10      DR Partner  completed
    Home Port: elb
    172.17.27.12      172.17.27.11      DR Auxiliary completed
24 entries were displayed.
cluster_A::>

```

9. Verify disk autoassignment and partitioning: `disk show -pool Pool1`

```
cluster_A::> disk show -pool Pool1
Disk              Usable      Disk      Container  Container
                  Size Shelf Bay  Type      Type      Name      Owner
-----
1.10.4            -          10    4 SAS      remote    -          node_B_2
1.10.13           -          10    13 SAS     remote    -          node_B_2
1.10.14           -          10    14 SAS     remote    -          node_B_1
1.10.15           -          10    15 SAS     remote    -          node_B_1
1.10.16           -          10    16 SAS     remote    -          node_B_1
1.10.18           -          10    18 SAS     remote    -          node_B_2
...
2.20.0            546.9GB   20    0 SAS      aggregate aggr0_rhal_a1 node_a_1
2.20.3            546.9GB   20    3 SAS      aggregate aggr0_rhal_a2 node_a_2
2.20.5            546.9GB   20    5 SAS      aggregate rhal_a1_aggr1 node_a_1
2.20.6            546.9GB   20    6 SAS      aggregate rhal_a1_aggr1 node_a_1
2.20.7            546.9GB   20    7 SAS      aggregate rhal_a2_aggr1 node_a_2
2.20.10           546.9GB   20    10 SAS     aggregate rhal_a1_aggr1 node_a_1
...
43 entries were displayed.

cluster_A::>
```

10. Mirror the root aggregates: `storage aggregate mirror -aggregate aggr0_node_A_3-IP`

Note: You must complete this step on each MetroCluster IP node.

```
cluster_A::> aggr mirror -aggregate aggr0_node_A_3-IP

Info: Disks would be added to aggregate "aggr0_node_A_3-IP" on node "node_A_3-IP"
in the following manner:

Second Plex

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

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-FC
  349.0GB   16.84GB   95% online    1 node_A_1-FC  raid_dp,
mirrored,
normal
aggr0_node_A_2-FC
  349.0GB   16.84GB   95% online    1 node_A_2-FC  raid_dp,
mirrored,
normal
aggr0_node_A_3-IP
  467.6GB   22.63GB   95% online    1 node_A_3-IP  raid_dp,
mirrored,
normal
aggr0_node_A_4-IP
  467.6GB   22.62GB   95% online    1 node_A_4-IP  raid_dp,
mirrored,
normal
aggr_data_a1
  1.02TB    1.01TB    1% online    1 node_A_1-FC  raid_dp,
mirrored,
normal
aggr_data_a2
  1.02TB    1.01TB    1% online    1 node_A_2-FC  raid_dp,
mirrored,
```

Finalizing the addition of the MetroCluster IP nodes

You must incorporate the new DR group into the MetroCluster configuration and create mirrored data aggregates on the new nodes.

Procedure

1. 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`

Note: 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)
Physical                               Usable
Size      Position  Disk                               Type      Size
-----
-----
-          dparity   5.10.15                           SAS        -
-          parity    5.10.16                           SAS        -
-          data      5.10.17                           SAS        546.9GB
547.1GB   data      5.10.18                           SAS        546.9GB
558.9GB   data      5.10.19                           SAS        546.9GB
558.9GB

Second Plex
      RAID Group rg0, 5 disks (block checksum, raid_dp)
Physical                               Usable
Size      Position  Disk                               Type      Size
-----
-----
-          dparity   4.20.17                           SAS        -
-          parity    4.20.14                           SAS        -
-          data      4.20.18                           SAS        546.9GB
547.1GB   data      4.20.19                           SAS        546.9GB
547.1GB   data      4.20.16                           SAS        546.9GB
547.1GB

Aggregate capacity available for volume use would be 1.37TB.

```

```
Do you want to continue? {y|n}: y
[Job 440] Job succeeded: DONE
```

```
cluster_A::>
```

2. Configure the MetroCluster to implement the changes: metrocluster configure

```
cluster_A::*> metrocluster configure
[Job 439] Job succeeded: Configure is successful.
cluster_A::*>
```

3. Verify that the nodes are added to their DR group: metrocluster node show

```
cluster_A::*> metrocluster node show
```

```

DR
Group Cluster Node Configuration State DR Mirroring Mode
-----
1 cluster_A
  node-A-1-FC configured enabled normal
  node-A-2-FC configured enabled normal
  Cluster-B
  node-B-1-FC configured enabled normal
  node-B-2-FC configured enabled normal
2 cluster_A
  node-A-3-IP configured enabled normal
  node-A-4-IP configured enabled normal
  Cluster-B
  node-B-3-IP configured enabled normal
  node-B-4-IP 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:

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

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

```

cluster_A::> volume show
Vserver Volume Aggregate State Type Size Available Used%
-----
...
cluster_A MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_A
  aggr_b1 - RW - - -
cluster_A MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_B
  aggr_b2 - RW - - -
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_A
  aggr_a1 online RW 10GB 9.50GB 0%
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
  aggr_a2 online RW 10GB 9.50GB 0%
...
11 entries were displayed.mple

```

b) Set the advanced privilege level: `set -privilege advanced`

c) Move the MDV volumes, one at a time: `volume move start -volume mdv-volume -destination-aggregate aggr-on-new-node -vserver vserver-name`

The following example shows the command and output for moving

MDV_CRS_d6b0b313ff5611e9837100a098544e51_A to aggregate data_a3 on node_A_3.

```

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

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

```

d) Use the `volume show` command to check that the MDV volume has been successfully moved:

`volume show mdv-name`

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

```

cluster_A::> vol show MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
Vserver Volume Aggregate State Type Size
Available Used%
-----
cluster_A MDV_CRS_d6b0b313ff5611e9837100a098544e51_B
  aggr_a2 online RW 10GB
9.50GB 0%

```


e) Return to admin mode: `set -privilege admin`

Moving the data to the new drive shelves

During the transition, you move data from the drive shelves in the MetroCluster FC configuration to the new MetroCluster IP configuration.

Procedure

- To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.
 - Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=end`
 - Repeat the command on the partner cluster.
- Move the data volumes to aggregates on the new controllers, one volume at a time.
Use the following section of the *Controller Upgrade Express Guide*.
Creating an aggregate and moving volumes to the new nodes
- Create SAN LIFs on the recently added nodes.
Use the following section of the *Cluster Expansion Express Guide*.
Updating LUN paths for the new nodes
- Check if there are any node locked licenses on the FC nodes, if there are, they need to be added to the newly added nodes.
Use the following section of the *Cluster Expansion Express Guide*.
Adding node-locked licenses
- Migrate the data LIFs.
Use the following section of the *Controller Upgrade Express Guide* but do **not** perform the last two steps to migrate cluster management LIFs.
Moving non-SAN data LIFs and cluster management LIFs to the new nodes

Removing the MetroCluster FC controllers

You must perform clean-up tasks and remove the old controller modules from the MetroCluster configuration.

Procedure

- To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway.
 - Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours 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`
 - Repeat the command on the partner cluster.
- Identify the aggregates hosted on the MetroCluster FC configuration that need to be deleted.
In this example the following data aggregates are hosted by the MetroCluster FC cluster_B and need to be deleted: `aggr_data_a1` and `aggr_data_a2`.

Note: You need to perform the steps to identify, offline and delete the data aggregates on both the clusters. The example is for one cluster only.

```
cluster_B::> aggr show
```

Aggregate	Size	Available	Used%	State	#Vols	Nodes	RAID Status
aggr0_node_A_1-FC	349.0GB	16.83GB	95%	online	1	node_A_1-FC	raid_dp, mirrored, normal
aggr0_node_A_2-FC	349.0GB	16.83GB	95%	online	1	node_A_2-FC	raid_dp,

```

mirrored,
normal
aggr0_node_A_3-IP
  467.6GB  22.63GB  95% online    1 node_A_3-IP    raid_dp,
mirrored,
normal
aggr0_node_A_3-IP
  467.6GB  22.62GB  95% online    1 node_A_4-IP    raid_dp,
mirrored,
normal
aggr_data_a1
  1.02TB   1.02TB   0% online     0 node_A_1-FC    raid_dp,
mirrored,
normal
aggr_data_a2
  1.02TB   1.02TB   0% online     0 node_A_2-FC    raid_dp,
mirrored,
normal
aggr_data_a3
  1.37TB   1.35TB   1% online     3 node_A_3-IP    raid_dp,
mirrored,
normal
aggr_data_a4
  1.25TB   1.24TB   1% online     2 node_A_4-IP    raid_dp,
mirrored,
normal
8 entries were displayed.
cluster_B::>

```

3. Check if the data aggregates on the FC nodes have any MDV_aud volumes, and delete them prior to deleting the aggregates.

You must delete the MDV_aud volumes as they cannot be moved.

4. Take each of the aggregates offline, and then delete them:

- a) Take the aggregate offline: `storage aggregate offline -aggregate aggregate-name`
The following example shows the aggregate `node_B_1_aggr0` being taken offline:

```

cluster_B::> storage aggregate offline -aggregate node_B_1_aggr0
Aggregate offline successful on aggregate: node_B_1_aggr0

```

- b) Delete the aggregate: `storage aggregate delete -aggregate aggregate-name`
You can destroy the plex when prompted.

The following example shows the aggregate `node_B_1_aggr0` being deleted.

```

cluster_B::> storage aggregate delete -aggregate node_B_1_aggr0
Warning: Are you sure you want to destroy aggregate "node_B_1_aggr0"? {y|n}: y
[Job 123] Job succeeded: DONE
cluster_B::>

```

5. Identify the MetroCluster FC DR group that need to be removed.

In the following example the MetroCluster FC nodes are in DR Group '1', and this is the DR group that need to be removed.

```

cluster_B::> metrocluster node show
DR
Group Cluster Node          Configuration State  DR
-----
1      cluster_A
      node_A_1-FC             configured enabled normal
      node_A_2-FC             configured enabled normal
      cluster_B
      node_B_1-FC             configured enabled normal
      node_B_2-FC             configured enabled normal
2      cluster_A
      node_A_3-IP             configured enabled normal
      node_A_4-IP             configured enabled normal
      cluster_B
      node_B_3-IP             configured enabled normal
      node_B_3-IP             configured enabled normal
8 entries were displayed.
cluster_B::>

```

6. Move the cluster management LIF from a MetroCluster FC node to a MetroCluster IP node:

```

cluster_B::> network interface migrate -vserver svm-name -lif
cluster_mgmt -destination-node node-in-metrocluster-ip-dr-group -
destination-port available-port

```

7. Change the home node and home port of the cluster management LIF: `cluster_B::> network interface modify -vserver svm-name -lif cluster_mgmt -service-policy default-management -home-node node-in-metrocluster-ip-dr-group -home-port lif-port`

8. Move epsilon from a MetroCluster FC node to a MetroCluster IP node:

- a) Identify which node currently has epsilon:
- `cluster show -fields epsilon`

```
cluster_B::> cluster show -fields epsilon
node          epsilon
-----
node_A_1-FC   true
node_A_2-FC   false
node_A_1-IP   false
node_A_2-IP   false
4 entries were displayed.
```

- b) Set epsilon to false on the MetroCluster FC node (node_A_1-FC):
- `cluster modify -node fc-node -epsilon false`

- c) Set epsilon to true on the MetroCluster IP node (node_A_1-IP):
- `cluster modify -node ip-node -epsilon true`

- d) Verify that epsilon has moved to the correct node:
- `cluster show -fields epsilon`

```
cluster_B::> cluster show -fields epsilon
node          epsilon
-----
node_A_1-FC   false
node_A_2-FC   false
node_A_1-IP   true
node_A_2-IP   false
4 entries were displayed.
```

9. On each cluster, remove the DR group containing the old nodes from the MetroCluster FC configuration.

You must perform this step on both clusters, one at a time.

```
cluster_B::> metrocluster remove-dr-group -dr-group-id 1

Warning: Nodes in the DR group that are removed from the MetroCluster
configuration will lose their disaster recovery protection.

Local nodes "node_A_1-FC, node_A_2-FC" will be removed from the
MetroCluster configuration. You must repeat the operation on the
partner cluster "cluster_B" to remove the remote nodes in the DR group.
Do you want to continue? {y|n}: y

Info: The following preparation steps must be completed on the local and partner
clusters before removing a DR group.

1. Move all data volumes to another DR group.
2. Move all MDV_CRS metadata volumes to another DR group.
3. Delete all MDV_aud metadata volumes that may exist in the DR group to
be removed.
4. Delete all data aggregates in the DR group to be removed. Root
aggregates are not deleted.
5. Migrate all data LIFs to home nodes in another DR group.
6. Migrate the cluster management LIF to a home node in another DR group.
Node management and inter-cluster LIFs are not migrated.
7. Transfer epsilon to a node in another DR group.

The command is vetoed if the preparation steps are not completed on the
local and partner clusters.
Do you want to continue? {y|n}: y
[Job 513] Job succeeded: Remove DR Group is successful.

cluster_B::>
```

10. Verify that the nodes are ready to be removed from the clusters.

You must perform this step on both clusters.

Note: At this point, the `metrocluster node show` command only shows the local MetroCluster FC nodes and no longer shows the nodes that are part of the partner cluster.

```
cluster_B::> metrocluster node show

DR          Configuration  DR
Group Cluster Node      State          Mirroring Mode
-----
1          cluster_A
           node_A_1-FC    ready to configure  -
           node_A_2-FC    ready to configure  -
2          cluster_A
           node_A_3-IP    configured          enabled  normal
           node_A_4-IP    configured          enabled  normal
           cluster_B
           node_B_3-IP    configured          enabled  normal
           node_B_4-IP    configured          enabled  normal
6 entries were displayed.

cluster_B::>
```

11. Disable storage failover for the MetroCluster FC nodes.

You must perform this step on each node.

```
cluster_A:>> storage failover modify -node node_A_1-FC -enabled false
cluster_A:>> storage failover modify -node node_A_2-FC -enabled false
cluster_A:>>
```

12. Unjoin the MetroCluster FC nodes from the clusters: `cluster unjoin -node node-name`

You must perform this step on each node.

```
cluster_A:>> cluster unjoin -node node_A_1-FC

Warning: This command will remove node "node_A_1-FC" from the cluster. You must
         remove the failover partner as well. After the node is removed, erase
         its configuration and initialize all disks by using the "Clean
         configuration and initialize all disks (4)" option from the boot menu.
Do you want to continue? {y|n}: y
[Job 553] Job is queued: Cluster remove-node of Node:node_A_1-FC with UUID:6c87de7e-ff54-11e9-8371
[Job 553] Checking prerequisites
[Job 553] Cleaning cluster database
[Job 553] Job succeeded: Node remove succeeded
If applicable, also remove the node's HA partner, and then clean its configuration and initialize all
disks with the boot menu.
Run "debug vreport show" to address remaining aggregate or volume issues.

cluster_B:>>
```

13. Power down the MetroCluster FC controller modules and storage shelves.

14. Disconnect and remove the MetroCluster FC controller modules and storage shelves.

Completing the transition

To complete the transition you must verify the operation of the new MetroCluster IP configuration.

Procedure

1. Verify the MetroCluster IP configuration.

You must perform this step on each cluster.

The following example shows the output for cluster_A.

```
cluster_A:>> cluster show
Node          Health  Eligibility  Epsilon
-----
node_A_1-IP   true   true         true
node_A_2-IP   true   true         false
2 entries were displayed.

cluster_A:>>
```

The following example shows the output for cluster_B.

```
cluster_B:>> cluster show
Node          Health  Eligibility  Epsilon
-----
node_B_1-IP   true   true         true
node_B_2-IP   true   true         false
2 entries were displayed.

cluster_B:>>
```

2. Enable cluster HA and storage failover.

You must perform this step on each cluster.

3. Verify that cluster HA capability is enabled.

```
cluster_A:>> cluster ha show
High Availability Configured: true

cluster_A:>>
```

```
cluster_A:>> storage failover show
Node          Partner      Takeover
Possible     State Description
-----
node_A_1-IP   node_A_2-IP  true     Connected to node_A_2-IP
node_A_2-IP   node_A_1-IP  true     Connected to node_A_1-IP
2 entries were displayed.

cluster_A:>>
```

4. Disable MetroCluster transition mode.

a) Change to the advanced privilege level: `set -privilege advanced`

b) Disable transition mode: `metrocluster transition disable`

```
cluster_A:*> metrocluster transition disable

cluster_A:*>
```

a) Return to the admin privilege level: `set -privilege admin`

5. Verify that transition is disabled: `metrocluster transition show-mode`

You must perform these steps on both clusters.

```
cluster_A:> metrocluster transition show-mode
Transition Mode
-----
not-enabled
cluster_A:>
```

```
cluster_B:> metrocluster transition show-mode
Transition Mode
-----
not-enabled
cluster_B:>
```

Sending a custom AutoSupport message after maintenance

After completing the transition, you should send an AutoSupport message indicating the end of maintenance, so automatic case creation can resume.

Procedure

To resume automatic support case generation, send an Autosupport message to indicate that the maintenance is complete.

- a) Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=end`
- b) Repeat the command on the partner cluster.

Restoring Tiebreaker or Mediator monitoring

After completing the transition of the MetroCluster configuration, you can resume monitoring with the Tiebreaker or Mediator utility.

Procedure

Use the appropriate procedure for your configuration.

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations in the Tiebreaker Software 1.21 Installation and Configuration Guide
Mediator	Configuring the ONTAP Mediator service from a MetroCluster IP configuration in the MetroCluster IP Installation and Configuration Guide

Moving SAN FC hosts from MetroCluster FC to MetroCluster IP nodes

After you transition your MetroCluster nodes from FC to IP, you might need to move your SAN FC host connections to the new nodes.

Procedure

1. Set up new FC interfaces (LIFs) on MetroCluster IP nodes:
 - a) If required, on MetroCluster IP nodes, modify FCP ports to be used for client connectivity to FCP target personality.
 - b) Create FCP LIFs/interfaces on IP nodes for all SAN vservers. Optionally verify that the WWPNs from newly created FCP LIFs are logged into the SAN switch
2. Updating SAN zoning configuration for newly added FCP LIFs on MetroCluster IP nodes.

In order to facilitate moving of volumes that contain LUNs actively serving data to FCP SAN clients, update existing FC switch zones to allow FC SAN clients to access to LUNs on MetroCluster IP nodes.

- a) On the FC SAN switch (Cisco or Brocade), add the WWPNs of newly added FC SAN LIFs to the zone.
- b) Update, save and commit the zoning changes.
- c) From the client, check for FCP initiator logins to the new SAN LIFs on the MetroCluster IP nodes:

```
sanlun lun show -p
```

At this time, client can see all FC interfaces (on MetroCluster FC and MetroCluster IP nodes) in ONTAP and is logged in to all those interfaces. LUNs and volumes are still physically hosted on FC nodes.

Because LUNs are reported only on FC node interfaces, the client shows only paths over FC nodes. This can be seen in the output of the `sanlun lun show -p` and `multipath -ll -d` commands.

```
[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdk host3 iscsi_lf_n2_p1_
up secondary sdh host2 iscsi_lf_n1_p1_

[root@scspr1789621001 ~]# multipath -ll -d
3600a098038304646513f4f674e52774b dm-5 NETAPP ,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|+- policy='service-time 0' prio=50 status=active
|`- 3:0:0:4 sdk 8:160 active ready running
`+- policy='service-time 0' prio=10 status=enabled
`- 2:0:0:4 sdh 8:112 active ready running
```

3. Change the reporting nodes to add IP nodes
 - a) List reporting nodes for LUNs on the SVM: `lun mapping show -vserver svm-name -fields reporting-nodes -ostype linux`

Reporting nodes shown are local nodes as LUNs are physically on FC nodes `sti8200mcc-htp-001` and `sti8200mcc-htp-002`.

```
sti8200mcchtp001htp_siteA:> lun mapping show -vserver vsa_1 -fields reporting-nodes -ostype linux
```

```

vserver  path                                igroup          reporting-nodes
-----  -
vsa_1    /vol/vsa_1_vol1/lun_linux_2           igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol1/lun_linux_3           igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol2/lun_linux_4           igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol3/lun_linux_7           igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol4/lun_linux_8           igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol4/lun_linux_9           igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol6/lun_linux_12          igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol6/lun_linux_13          igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol7/lun_linux_14          igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol8/lun_linux_17          igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol9/lun_linux_18          igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
vsa_1    /vol/vsa_1_vol9/lun_linux_19          igroup_linux    sti8200mcc-htp-001,sti8200mcc-htp-002
12 entries were displayed.

```

- b) Add reporting nodes to include MetroCluster IP nodes.

```

sti8200mcchtp001htp_siteA::> lun mapping add-reporting-nodes -
vserver vsa_1 -path /vol/vsa_1_vol*/lun_linux_* -nodes sti8200mccip-
htp-005,sti8200mccip-htp-006 -igroup igroup_linux

```

12 entries were acted on.

- c) List reporting nodes and verify the presence of the new nodes:

```

sti8200mcchtp001htp_siteA::> lun mapping show -vserver vsa_1 -fields reporting-nodes -ostype linux
vserver path                                igroup          reporting-nodes
-----  -
vsa_1    /vol/vsa_1_vol1/lun_linux_2           igroup_linux    sti8200mcc-htp-001,sti8200mcc-
htp-002,sti8200mccip-htp-005,sti8200mccip-htp-006
vsa_1    /vol/vsa_1_vol1/lun_linux_3           igroup_linux    sti8200mcc-htp-001,sti8200mcc-
htp-002,sti8200mccip-htp-005,sti8200mccip-htp-006
vsa_1    /vol/vsa_1_vol2/lun_linux_4           igroup_linux    sti8200mcc-htp-001,sti8200mcc-
htp-002,sti8200mccip-htp-005,sti8200mccip-htp-006
vsa_1    /vol/vsa_1_vol3/lun_linux_7           igroup_linux    sti8200mcc-htp-001,sti8200mcc-
htp-002,sti8200mccip-htp-005,sti8200mccip-htp-006
...

```

12 entries were displayed.

- d) Rescan the scsi bus on the host to discover the newly added paths: `/usr/bin/rescan-scsi-bus.sh -a`

```

[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -a
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning for device 2 0 0 0 ...
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.

```

- e) Display the newly added paths: `sanlun lun show -p`

Each LUN will have four paths.

```

[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF

```

```

-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
up secondary sdag host4 iscsi_lf__n4_p1_
up secondary sdah host5 iscsi_lf__n3_p1_

```

- f) On the controllers, move the volume/volumes containing LUNs from FC to IP nodes.

```

sti8200mcchtp001htp_siteA::> vol move start -vserver vsa_1 -volume
vsa_1_voll -destination-aggregate sti8200mccip_htp_005_aggr1
[Job 1877] Job is queued: Move "vsa_1_voll" in Vserver "vsa_1" to
aggregate "sti8200mccip_htp_005_aggr1". Use the "volume move show -
vserver vsa_1 -volume vsa_1_voll"
command to view the status of this operation.
sti8200mcchtp001htp_siteA::> volume move show
Vserver      Volume      State      Move Phase      Percent-Complete Time-To-
Complete
-----
vsa_1        vsa_1_voll  healthy   initializing
--

```

- g) When the volume move is completed, check volume and LUN status to confirm that the volume and/or lun is not offline (optional step)
- h) On the FC SAN client, display the LUN information: `sanlun lun show -p`

The iSCSI interfaces on the IP nodes on which the LUN now resides are updated as primary paths. If the primary path is not updated after vol move, run `"/usr/bin/rescan-scsi-bus.sh -a"` or simply wait for multipath rescanning to take place.

The primary path in the following example is the LIF on MetroCluster IP node.

```

[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdag host4 iscsi_lf__n4_p1_
up secondary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_
up secondary sdah host5 iscsi_lf__n3_p1_

```

All LUNs for the SVM are now on the MetroCluster IP nodes.

4. Remove the reporting nodes and re-scan paths from client.

- a) Remove the remote reporting nodes (the MetroCluster FC nodes) for the linux LUNs: `lun mapping remove-reporting-nodes -vserver vsa_1 -path * -igroup igroup_linux -remote-nodes true`

```

sti8200mcchtp001htp_siteA::> lun mapping remove-reporting-nodes -
vserver vsa_1 -path * -igroup igroup_linux -remote-nodes true
12 entries were acted on.

```

- b) Check reporting nodes for the LUNs: `lun mapping show -vserver vsa_1 -fields reporting-nodes -ostype linux`

```

sti8200mcchtp001htp_siteA::> lun mapping show -vserver vsa_1 -fields
reporting-nodes -ostype linux

```



```
vserver path igroup reporting-nodes
-----
vsa_1 /vol/vsa_1_vol1/lun_linux_2 igroup_linux sti8200mccip-
htp-005,sti8200mccip-htp-006
vsa_1 /vol/vsa_1_vol1/lun_linux_3 igroup_linux sti8200mccip-
htp-005,sti8200mccip-htp-006
vsa_1 /vol/vsa_1_vol2/lun_linux_4 igroup_linux sti8200mccip-
htp-005,sti8200mccip-htp-006
...

12 entries were displayed.
```

- c) Rescan the scsi bus on the client: `/usr/bin/rescan-scsi-bus.sh -r`
 The paths from the MetroCluster FC nodes are removed:

```
[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
sg0 changed: LU not available (PQual 1)
REM: Host: scsi2 Channel: 00 Id: 00 Lun: 00
DEL: Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
sg2 changed: LU not available (PQual 1)
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
24 device(s) removed.
[2:0:0:0]
[2:0:0:1]
...
```

- d) Verify that only paths from the MetroCluster IP nodes are visible from the host: `sanlun lun show -p`
 e) If required, remove iSCSI LIFs from the MetroCluster FC nodes.
 This should be done if there are no other LUNs on the nodes mapped to other clients.

Moving Linux iSCSI hosts from MetroCluster FC to MetroCluster IP nodes

After you transition your MetroCluster nodes from FC to IP, you might need to move your iSCSI host connections to the new nodes.

About this task

In this procedure IPv4 interfaces are created.

Setting up new iSCSI connections

To move iSCSI connections, you must set up new iSCSI connections to the MetroCluster IP nodes.

About this task

Procedure

1. Create iSCSI interfaces on the MetroCluster IP nodes and check ping connectivity from the iSCSI clients to the new IP interfaces on the MetroCluster IP nodes.

Creating network interfaces

All iSCSI interfaces from the SVM should be reachable by iSCSI client.

2. On the iSCSI host or client, identify the existing iSCSI connections from the host to the MetroCluster IP node: `iscsiadm -m session`

```
[root@scspr1789621001 ~]# iscsiadm -m session
tcp: [1] 10.230.68.236:3260,1156
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-
flash)
tcp: [2] 10.230.68.237:3260,1158
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-
flash)
```

3. Verify the connections from the MetroCluster IP node: `iscsi session show -vserver svm-name`

```
node_A_1-IP:*> iscsi session show -vserver vsa_1

Tpgroup Initiator Initiator

Vserver Name TSIH Name ISID Alias
-----
vsa_1 iscsi_lf_n1_p1_4
iqn.2020-01.com.fujitsu.englab.gdl:scspr1789621001 00:02:3d:00:00:01
scspr1789621001.gdl.englab.fujitsu.com
vsa_1 iscsi_lf_n2_p1_4
iqn.2020-01.com.fujitsu.englab.gdl:scspr1789621001 00:02:3d:00:00:02
scspr1789621001.gdl.englab.fujitsu.com

2 entries were displayed.
```

4. List the iscsi interfaces in ONTAP for the SVM that contains the interfaces: `iscsi interface show -vserver svm-name`

```
sti8200mcchtp001hwp_siteA:*> iscsi interface show -vserver vsa_1

Logical Status Curr Curr

Vserver Interface TPGT Admin/Oper IP Address Node Port Enabled
```

```
-----
-----
vsa_1 iscsi_lf__n1_p1_ 1156 up/up 10.230.68.236 sti8200mcc-htp-001 e0g
true
vsa_1 iscsi_lf__n1_p2_ 1157 up/up fd20:8b1e:b255:805e::78c9 sti8200mcc-
htp-001 e0h true
vsa_1 iscsi_lf__n2_p1_ 1158 up/up 10.230.68.237 sti8200mcc-htp-002 e0g
true
vsa_1 iscsi_lf__n2_p2_ 1159 up/up fd20:8b1e:b255:805e::78ca sti8200mcc-
htp-002 e0h true
vsa_1 iscsi_lf__n3_p1_ 1183 up/up 10.226.43.134 sti8200mccip-htp-005 e0c
true
vsa_1 iscsi_lf__n4_p1_ 1188 up/up 10.226.43.142 sti8200mccip-htp-006 e0c
true
```

6 entries were displayed.

5. On the iSCSI client, run discovery on any one of the iSCSI IP addresses on the SVM to discover the new targets: `iscsiadm -m discovery -t sendtargets -p iscsi-ip-address`

Discovery can be run on any IP address of the SVM, including non-iSCSI interfaces.

```
[root@scspr1789621001 ~]# iscsiadm -m discovery -t sendtargets -p
10.230.68.236:3260
```

```
10.230.68.236:3260,1156
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.142:3260,1188
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.226.43.134:3260,1183
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
10.230.68.237:3260,1158
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6
```

6. On the iSCSI client, login to all the discovered addresses: `iscsiadm -m node -L all -T node-address -p portal-address -l`

```
[root@scspr1789621001 ~]# iscsiadm -m node -L all -T
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 -p
10.230.68.236:3260 -l
```

```
Logging in to [iface: default, target:
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6,
portal: 10.226.43.142,3260] (multiple)
Logging in to [iface: default, target:
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6,
portal: 10.226.43.134,3260] (multiple)
Login to [iface: default, target:
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6,
portal: 10.226.43.142,3260] successful.
Login to [iface: default, target:
iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6,
portal: 10.226.43.134,3260] successful.
```

7. On the iSCSI client, verify the login and connections: `iscsiadm -m session`

```
[root@scspr1789621001 ~]# iscsiadm -m session

tcp: [1] 10.230.68.236:3260,1156 iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-
flash)
tcp: [2] 10.230.68.237:3260,1158 iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-
flash)
tcp: [3] 10.226.43.142:3260,1188 iqn.1992-08.com.fujitsu:sn.58d7f6df2cc611eaa9c500a098a71638:vs.6 (non-
flash)
```

8. From the MetroCluster node, verify the login and connection with the client: `iscsi initiator show -vserver vsa_1`

```
sti8200mcchtp001htp_siteA::*> iscsi initiator show -vserver vsa_1
```

```
Tpgroup Initiator

Vserver Name          TSIH Name          ISID          Igroup Name
-----
```

```
vsa_1 iscsi_lf__n1_p1_ 4 iqn.2020-01.com.fujitsu.enlab.gdl:scspr1789621001 00:02:3d:00:00:01
igroup_linux
vsa_1 iscsi_lf__n2_p1_ 4 iqn.2020-01.com.fujitsu.enlab.gdl:scspr1789621001 00:02:3d:00:00:02
igroup_linux
vsa_1 iscsi_lf__n3_p1_ 1 iqn.2020-01.com.fujitsu.enlab.gdl:scspr1789621001 00:02:3d:00:00:04
igroup_linux
vsa_1 iscsi_lf__n4_p1_ 1 iqn.2020-01.com.fujitsu.enlab.gdl:scspr1789621001 00:02:3d:00:00:03
igroup_linux

4 entries were displayed.
```

Results

At the end of this tasks, the client can see all iSCSI interfaces (on the MetroCluster FC and MetroCluster IP nodes) and is logged in to all those interfaces.

LUNs and volumes are still physically hosted on FC nodes. Because LUNs are reported only on MetroCluster FC node interfaces, the client will show only paths over MetroCluster FC nodes. This can be seen in `sanlun lun show -p` and `multipath -ll -d` command outputs. The next step is to add IP nodes as reporting nodes.

```
[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----

host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdk host3 iscsi_lf__n2_p1_
up secondary sdh host2 iscsi_lf__n1_p1_

[root@scspr1789621001 ~]# multipath -ll -d
3600a098038304646513f4f674e52774b dm-5 NETAPP ,LUN C-Mode
size=2.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|--+ policy='service-time 0' prio=50 status=active
| `-- 3:0:0:4 sdk 8:160 active ready running
|--+ policy='service-time 0' prio=10 status=enabled
| `-- 2:0:0:4 sdh 8:112 active ready running
```

Adding the MetroCluster IP nodes as reporting nodes

After setting up the connections to the new MetroCluster IP nodes, you must add new reporting nodes.

Procedure

1. On the MetroCluster node, list reporting nodes for LUNs on the SVM: `lun mapping show -vserver vsa_1 -fields reporting-nodes -ostype linux`
The following reporting nodes are local nodes as LUNs are physically on FC nodes `node_A_1-FC` and `node_A_2-FC`.

```
node_A_1-IP::*> lun mapping show -vserver vsa_1 -fields reporting-nodes -
ostype linux

vserver path igroup reporting-nodes
-----
vsa_1 /vol/vsa_1_vol1/lun_linux_2 igroup_linux node_A_1-FC,node_A_2-FC
.
.
.
vsa_1 /vol/vsa_1_vol9/lun_linux_19 igroup_linux node_A_1-FC,node_A_2-FC
```

```
12 entries were displayed.
```

2. On the MetroCluster node, add reporting nodes: `lun mapping add-reporting-nodes -vserver svm-name -path /vol/vsa_1_vol*/lun_linux_* -nodes node1,node2 -igroup igroup_linux`

```
node_A_1-IP:~> lun mapping add-reporting-nodes -vserver vsa_1 -path /
vol/vsa_1_vol*/lun_linux_* -nodes node_A_1-IP,node_A_2-IP
-igroup igroup_linux
```

```
12 entries were acted on.
```

3. On the MetroCluster node, verify that the newly added nodes are present: `lun mapping show -vserver svm-name -fields reporting-nodes -ostype linux vserver path igroup reporting-nodes`

```
node_A_1-IP:~> lun mapping show -vserver vsa_1 -fields reporting-nodes -
ostype linux vserver path igroup reporting-nodes
-----
```

```
vsa_1 /vol/vsa_1_vol1/lun_linux_2 igroup_linux node_A_1-FC,node_A_2-
FC,node_A_1-IP,node_A_2-IP
vsa_1 /vol/vsa_1_vol1/lun_linux_3 igroup_linux node_A_1-FC,node_A_2-
FC,node_A_1-IP,node_A_2-IP.
.
.
.
```

```
12 entries were displayed.
```

4. On the host, issue the following command to rescan the scsi bus on the host and discover the newly added paths: `/usr/bin/rescan-scsi-bus.sh -a`

```
[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -a
Scanning SCSI subsystem for new devices
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
  Scanning for device 2 0 0 0 ...
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
  Vendor: NETAPP Model: LUN C-Mode Rev: 9800
  Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
0 device(s) removed.
```

5. On the host, issue the following command to list the newly added paths: `sanlun lun show -p`
Four paths are shown for each LUN.

```
[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
```

```
host vserver
path path /dev/ host vserver
state type node adapter LIF
```

```

-----
up primary sdk host3 iscsi_lf_n2_p1_
up secondary sdh host2 iscsi_lf_n1_p1_
up secondary sdag host4 iscsi_lf_n4_p1_
up secondary sdah host5 iscsi_lf_n3_p1_

```

6. On the MetroCluster, move the volume/volumes containing LUNs from FC to IP nodes.

```

node_A_1-IP:~*~> vol move start -vserver vsa_1 -volume vsa_1_voll -
destination-aggregate sti8200mccip_htp_005_aggr1
[Job 1877] Job is queued: Move "vsa_1_voll" in Vserver "vsa_1" to
aggregate "sti8200mccip_htp_005_aggr1". Use the "volume move show -
vserver
vsa_1 -volume vsa_1_voll" command to view the status of this operation.
node_A_1-IP:~*~> vol move show
Vserver Volume State Move Phase Percent-
Complete Time-To-Complete
-----
vsa_1 vsa_1_voll healthy initializing - -

```

7. When the volume move is completed, check volume and LUN status to confirm that the volume or LUN is not offline from the client: `sanlun lun show -p`
8. The iSCSI interfaces on the MetroCluster IP nodes where the LUN now resides are updated as primary paths. If the primary path is not updated after vol move, run `/usr/bin/rescan-scsi-bus.sh -a` or simply wait for multipath rescanning to take place.
- In the following example, the primary path is a LIF on the MetroCluster IP node.

```

[root@scspr1789621001 ~]# sanlun lun show -p
ONTAP Path: vsa_1:/vol/vsa_1_vol6/lun_linux_12
LUN: 4
LUN Size: 2g
Product: cDOT
Host Device: 3600a098038304646513f4f674e52774b
Multipath Policy: service-time 0
Multipath Provider: Native
-----
host vserver
path path /dev/ host vserver
state type node adapter LIF
-----
up primary sdag host4 iscsi_lf_n4_p1_
up secondary sdk host3 iscsi_lf_n2_p1_
up secondary sdh host2 iscsi_lf_n1_p1_
up secondary sdah host5 iscsi_lf_n3_p1_

```

Removing reporting nodes and rescanning paths

You must remove the reporting nodes and rescan the paths.

Procedure

- From the MetroCluster IP node, remove remote reporting nodes (the MetroCluster IP nodes) for the Linux LUNs: `lun mapping remove-reporting-nodes -vserver vsa_1 -path * -igroup igroup_linux -remote-nodes true`
In this case, the remote nodes are FC nodes.

```

node_A_1-IP:~*~> lun mapping remove-reporting-nodes -vserver vsa_1 -path *
-igroup igroup_linux -remote-nodes true

12 entries were acted on.

```

- From the MetroCluster IP node, check reporting nodes for the LUNs: `lun mapping show -vserver vsa_1 -fields reporting-nodes -ostype linux`

```
node_A_1-IP::~*> lun mapping show -vserver vsa_1 -fields reporting-nodes -
ostype linux

vserver  path                                     igroup      reporting-nodes
-----  -
vsa_1    /vol/vsa_1_vol1/lun_linux_2                igroup_linux node_A_1-IP,node_A_2-IP
vsa_1    /vol/vsa_1_vol1/lun_linux_3                igroup_linux node_A_1-IP,node_A_2-IP
vsa_1    /vol/vsa_1_vol2/lun_linux_4                group_linux  node_A_1-IP,node_A_2-IP
.
.
.

12 entries were displayed.
```

3. On the iSCSI host, rescan the scsi bus: `/usr/bin/rescan-scsi-bus.sh -r`
The paths that are removed are the paths from FC nodes.

```
[root@scspr1789621001 ~]# /usr/bin/rescan-scsi-bus.sh -r
Syncing file systems
Scanning SCSI subsystem for new devices and remove devices that have
disappeared
Scanning host 0 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 1 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
Scanning host 2 for SCSI target IDs 0 1 2 3 4 5 6 7, all LUNs
sg0 changed: LU not available (PQual 1)
REM: Host: scsi2 Channel: 00 Id: 00 Lun: 00
DEL: Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
sg2 changed: LU not available (PQual 1)
.
.
.
OLD: Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: NETAPP Model: LUN C-Mode Rev: 9800
Type: Direct-Access ANSI SCSI revision: 05
0 new or changed device(s) found.
0 remapped or resized device(s) found.
24 device(s) removed.
[2:0:0:0]
[2:0:0:1]
.
.
.
```

4. On the iSCSI host, verify that only paths from the MetroCluster IP nodes are visible: `sanlun lun show -pmultipath`

Expanding a two-node MetroCluster FC configuration to a four-node configuration

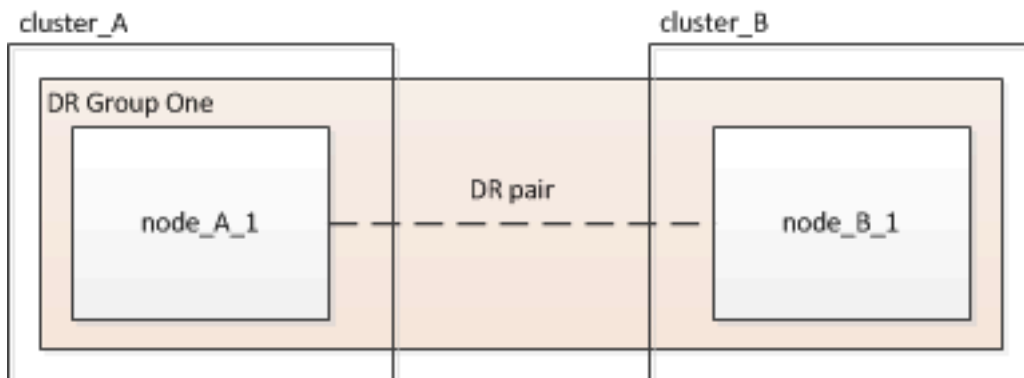
Expanding a two-node MetroCluster FC configuration to a four-node MetroCluster FC configuration involves adding a controller to each cluster to form an HA pair at each MetroCluster site, and then refreshing the MetroCluster FC configuration.

Before you begin


- The nodes must be running ONTAP 9 or later in a MetroCluster FC configuration.
- This procedure is not supported on earlier versions of ONTAP or in MetroCluster IP configurations.
- If the platforms in your two-node configuration are not supported in ONTAP 9.7 and you plan to upgrade to platforms supported in ONTAP 9.7 *and* expand to a four-node cluster, you must upgrade the platforms in the two-node configuration *before* expanding the MetroCluster FC configuration.
- The existing MetroCluster FC configuration must be healthy.
- 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.

About this task

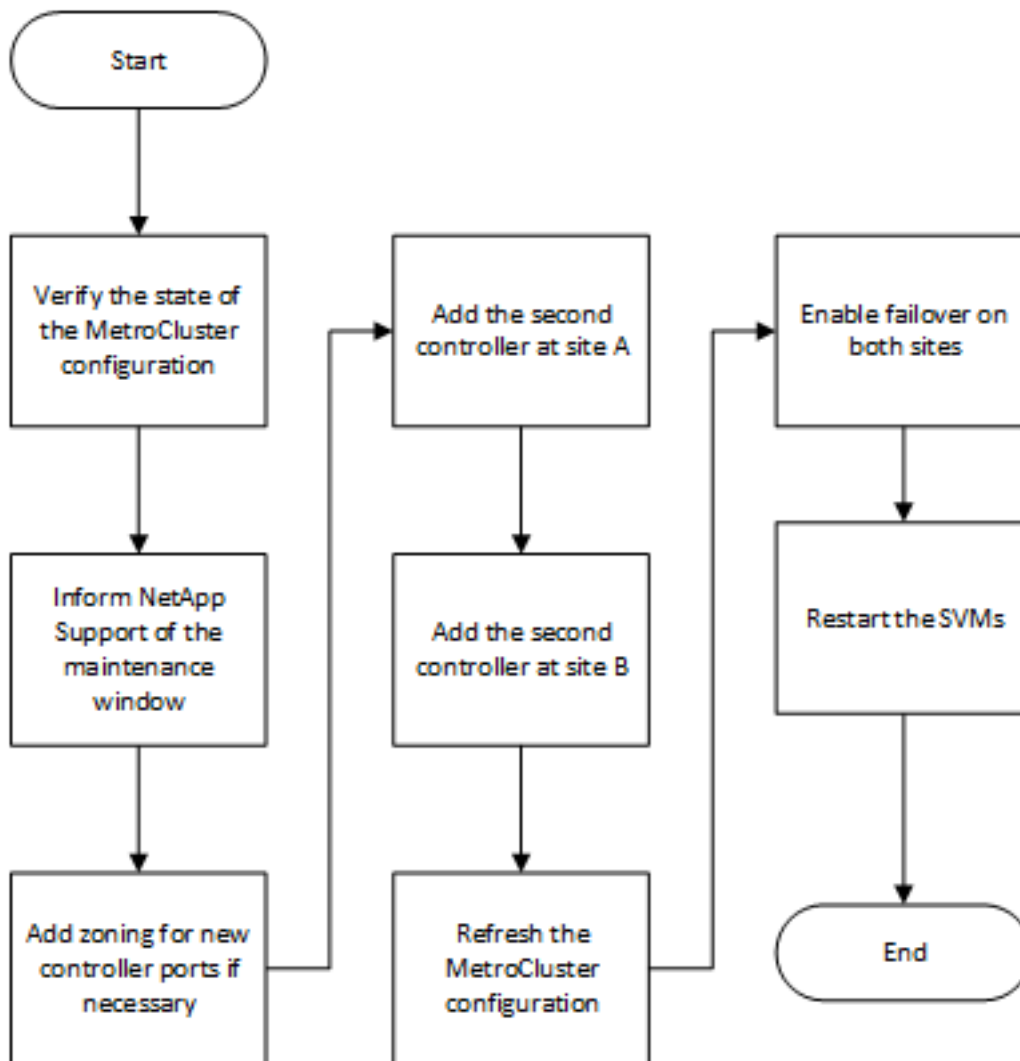
- This procedure applies only to MetroCluster FC configurations.
- This procedure is disruptive and takes approximately four hours to complete.
- Before performing this procedure, the MetroCluster FC configuration consists of two single-node clusters:



After completing this procedure, the MetroCluster FC configuration consists of two HA pairs, one at each site:



- Both sites must be expanded equally.
A MetroCluster configuration cannot consist of an uneven number of nodes.
- This procedure can take over an hour per site, with additional time for tasks such as initializing the disks and netbooting the new nodes.
The time to initialize the disks depends on the size of the disks.
- This procedure uses the following workflow:



Verifying the state of the MetroCluster configuration

You should identify the existing controllers and confirm the disaster recovery (DR) relationships between them, that the controllers are in normal mode, and that the aggregates are mirrored.

Procedure

1. Display the details of the nodes in the MetroCluster configuration from any node in the configuration: `metrocluster node show -fields node,dr-partner,dr-partner-systemid`

The following output shows that this MetroCluster configuration has a single DR group and one node in each cluster.

```

cluster_A::> metrocluster node show -fields node,dr-partner,dr-partner-systemid
dr-group-id  cluster      node          dr-partner    dr-partner-systemid
-----
1            cluster_A    controller_A_1 controller_B_1 536946192
1            cluster_B    controller_B_1 controller_A_1 536946165
2 entries were displayed.
  
```

2. Display the state of the MetroCluster configuration: `metrocluster show`

The following output shows that the existing nodes in the MetroCluster configuration are in normal mode:

```

cluster_A::> metrocluster show
Configuration: two-node-fabric

Cluster      Entry Name      State
-----
Local: cluster_A Configuration State configured
              Mode          normal
  
```

```
Remote: controller_B_1_siteB
AUSO Failure Domain      auso-on-cluster-disaster
Configuration State      configured
Mode                      normal
AUSO Failure Domain      auso-on-cluster-disaster
```

3. Check the state of the aggregates on each node in the MetroCluster configuration: `storage aggregate show`

The following output shows that the aggregates on cluster_A are online and mirrored:

```
cluster_A::> storage aggregate show

Aggregate          Size      Available Used%   State  #Vols  Nodes          RAID Status
-----
aggr0_controller_A_1_0  1.38TB   68.63GB   95%   online    1   controller_A_1  raid_dp,mirrored
controller_A_1_aggr1  4.15TB   4.14TB    0%   online    2   controller_A_1  raid_dp,mirrored
controller_A_1_aggr2  4.15TB   4.14TB    0%   online    1   controller_A_1  raid_dp,mirrored
3 entries were displayed.

cluster_A::>
```

Sending a custom AutoSupport message before adding nodes to the MetroCluster configuration

You should issue an AutoSupport message to notify Fujitsu support that maintenance is underway. Informing Fujitsu 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.

Procedure

1. Log in to the cluster at Site_A.
2. Invoke an AutoSupport message indicating the start of the maintenance: `system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours maintenance-window-in-hours` specifies the length of the maintenance window and can be a maximum of 72 hours. If you complete the maintenance before the time has elapsed, you can issue the `system node autosupport invoke -node * -type all -message MAINT=end` command to indicate that the maintenance period has ended.
3. Repeat this step on the partner site.

Zoning for the new controller ports when adding a controller module in a fabric-attached MetroCluster configuration

The FC switch zoning must accommodate the new controller connections. If you used the Fujitsu-supplied reference configuration files (RCFs) to configure your switches, the zoning is preconfigured and you do not need to make any changes.

About this task

If you manually configured your FC switches, you must ensure that the zoning is correct for the initiator connections from the new controller modules.

Adding a new controller module to each cluster

You must add a new controller module to each site, creating an HA pair in each site. This is a multistep process involving both hardware and software changes that must be performed in the proper order at each site.

Before you begin

- The new controller module must be received from Fujitsu as part of the upgrade kit.

You should verify that PCIe cards in the new controller module are compatible and supported by the new controller module.

- Your system must have an empty slot available for the new controller module when upgrading to a single-chassis HA pair (an HA pair in which both controller modules reside in the same chassis).

Note: This configuration is not supported on all systems. Platforms with single chassis configurations that are supported in ONTAP 9 are HX6100, AX4100.

- You must have rack space and cables for the new controller module when upgrading to a dual-chassis HA pair (an HA pair in which the controller modules reside in separate chassis).

Note: This configuration is not supported on all systems.

- You must connect each controller module to the management network through its e0a port or, if your system has one, you can connect to the e0M port as the management port.

About this task

- These tasks must be repeated at each site.
- The preexisting controller modules are referred to as the *existing* controller modules.

The examples in this procedure have the console prompt `existing_ctlr>`.

- The controller modules that are being added are referred to as the *new* controller modules; the examples in this procedure have the console prompt `new_ctlr>`.
- This task uses the following workflow:



Preparing for the upgrade

Before upgrading to an HA pair, you must verify that your system meets all requirements and that you have all of the necessary information.

Procedure

1. You need to identify unassigned disks or spare disks with available partitions that you can assign to the new controller module.
2. Based on the results of the previous step, perform either of the following:

If the result showed...	Then...
Not enough spare disks available for the new controller module on a system without root-data partitioning	<p>Contact Fujitsu support for more information.</p> <p>Complete the following substeps:</p> <ol style="list-style-type: none"> a. Determine where the aggregates for the existing node are located: <code>storage aggregate show</code> b. If disk ownership automatic assignment is on, turn it off: <code>storage disk option modify -node node_name -autoassign off</code> c. Remove ownership on disks that do not have aggregates on them: <code>storage disk removeowner disk_name</code> d. Repeat the previous step for as many disks as you need for the new node.

3. Verify that you have cables ready for the following connections:

- Cluster connections

If you are creating a two-node switchless cluster, you require two cables to connect the controller modules. Otherwise, you require a minimum of four cables, two for each controller module connection to the cluster-network switch. Other systems have defaults of either four or six cluster connections.

- HA interconnect connections, if the system is in a dual-chassis HA pair

4. Verify that you have a serial port console available for the controller modules.
5. Verify that your environment meets the site and system requirements.
6. Gather all of the IP addresses and other network parameters for the new controller module.

Clearing the configuration on a controller module

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

Procedure

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: `saveenvbye`
4. At the LOADER prompt, launch the boot menu: `boot_ontap menu`
5. At the boot menu prompt, clear the configuration: `wipeconfig`
Respond `yes` to the confirmation prompt.
The node reboots and the boot menu is displayed again.
6. At the boot menu, select option **5** to boot the system into Maintenance mode.
Respond `yes` to the confirmation prompt.

Preparing cluster ports on an existing controller module

Before installing a new controller module, you must configure cluster ports on the existing controller module so that the cluster ports can provide cluster communication with the new controller module.

About this task

If you are creating a two-node switchless cluster (with no cluster network switches), you must enable the switchless cluster networking mode.

For detailed information about port, LIF, and network configuration in ONTAP, see the *Network Management Guide*.

Procedure

1. Determine which ports should be used as the node's cluster ports.

The *Installation and Setup Instructions* for your platform on the Fujitsu Support Site contains information about the ports for cluster network connections.

2. For each cluster port, identify the port roles: `network port show`

In the following example, ports e0a, e0b, e0c, and e0d must be changed to cluster ports:

```
cluster_A::> network port show

Node: controller_A_1
Speed(Mbps) Health
Port    IPspace      Broadcast Domain Link  MTU    Admin/Oper  Status
-----
e0M     Default      mgmt_bd_1500  up    1500    auto/1000  healthy
e0a     Default      Default        up    1500    auto/10000 healthy
e0b     Default      Default        up    1500    auto/10000 healthy
e0c     Default      Default        up    1500    auto/10000 healthy
e0d     Default      Default        up    1500    auto/10000 healthy
e0i     Default      Default        down  1500    auto/10    -
e0j     Default      Default        down  1500    auto/10    -
e0k     Default      Default        down  1500    auto/10    -
e0l     Default      Default        down  1500    auto/10    -
e2a     Default      Default        up    1500    auto/10000 healthy
e2b     Default      Default        up    1500    auto/10000 healthy
e4a     Default      Default        up    1500    auto/10000 healthy
e4b     Default      Default        up    1500    auto/10000 healthy
13 entries were displayed.
```

3. For any data LIF that is using a cluster port as the home-port or current-port, modify the LIF to use a data port as its home-port: `network interface modify`

The following example changes the home port of a data LIF to a data port:

```
cluster1::> network interface modify -lif datalif1 -vserver vs1 -home-port e1b
```

4. For each LIF that you modified, revert the LIF to its new home port: `network interface revert`

The following example reverts the LIF datalif1 to its new home port e1b:

```
cluster1::> network interface revert -lif datalif1 -vserver vs1
```

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

- a) Delete VLAN ports: `network port vlan delete -node node-name -vlan-name portid-vlandid`
For example:

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

- b) Remove physical ports from the interface groups: `network port ifgrp remove-port -node node-name -ifgrp interface-group-name -port portid`
For example:

```
network port ifgrp remove-port -node node1 -ifgrp ala -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`

6. Verify that the port roles have changed: `network port show`

The following example shows that ports e0a, e0b, e0c, and e0d are now cluster ports:

```
Node: controller_A_1
Speed(Mbps) Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
-----
e0M Default mgmt_bd_1500 up 1500 auto/1000 healthy
e0a Cluster Cluster up 9000 auto/10000 healthy
e0b Cluster Cluster up 9000 auto/10000 healthy
e0c Cluster Cluster up 9000 auto/10000 healthy
e0d Cluster Cluster up 9000 auto/10000 healthy
e0i Default Default down 1500 auto/10 -
e0j Default Default down 1500 auto/10 -
e0k Default Default down 1500 auto/10 -
e0l Default Default down 1500 auto/10 -
e2a Default Default up 1500 auto/10000 healthy
e2b Default Default up 1500 auto/10000 healthy
e4a Default Default up 1500 auto/10000 healthy
e4b Default Default up 1500 auto/10000 healthy
13 entries were displayed.
```

7. If your system is part of a switched cluster, create cluster LIFs on the cluster ports: `network interface create`

The following example creates a cluster LIF on one of the node's cluster ports. The `-auto` parameter configures the LIF to use a link-local IP address.

```
cluster1::> network interface create -vserver Cluster -lif clus1 -role
cluster -home-node node0 -home-port e1a -auto true
```

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

- a) Change to the advanced privilege level from either node: `set -privilege advanced`
You can respond `y` when prompted whether you want to continue into advanced mode. The advanced mode prompt appears (`*>`).
- b) Enable the switchless cluster networking mode: `network options switchless-cluster modify -enabled true`
- c) Return to the admin privilege level: `set -privilege admin`

Important: Cluster interface creation for the existing node in a two-node switchless cluster system is completed after cluster setup is completed through a netboot on the new controller module.

[Network Management Guide](#)

Preparing the netboot server to download the image

When you are ready to prepare the netboot server, you must download the correct ONTAP netboot image from the DVD included in the Product to the netboot server and note the IP address.

Before you begin

- You must be able to access an HTTP server from the system before and after adding the new controller module.
- You must have access to the DVD included in the Product to download the necessary system files for your platform and your version of ONTAP.
- Both controller modules in the HA pair must run the same version of ONTAP.

Procedure

1. Download the appropriate ONTAP software from the software download section of the Fujitsu Support Site and store the `<ontap_version>_image.tgz` file on a web-accessible directory.
The `<ontap_version>_image.tgz` file is used for performing a netboot of your system.
2. Change to the web-accessible directory and verify that the files you need are available.

For...	Then...
All other systems	Your directory listing should contain the following file:

For...	Then...
	<code><ontap_version>_image.tgz</code>
	Note: There is no need to extract the file contents.

- Determine the IP address of the existing controller module.
This address is referred to later in this procedure as *ip-address-of-existing controller*.
- Ping *ip-address-of-existing controller* to verify that the IP address is reachable.

Setting the HA mode on the existing controller module

You must use the `storage failover modify` command to set the mode on the existing controller module. The mode value is enabled later, after you reboot the controller module.

Procedure

Set the mode to HA: `storage failover modify -mode ha -node existing_node_name`

Shutting down the existing controller module

You must perform a clean shutdown of the existing controller module to verify that all of the data has been written to disk. You must also disconnect the power supplies.

Procedure

- Halt the node from the existing controller module prompt: `halt local -inhibit-takeover true`
If you are prompted to continue the halt procedure, enter `y` when prompted, and then wait until the system stops at the `LOADER` prompt.



Attention: You must perform a clean system shutdown before replacing the system components to avoid losing unwritten data in the NVRAM or NVMEM.

This LED blinks if there is unwritten data in the NVRAM. If this LED is flashing amber after you enter the `halt` command, you need to reboot your system and try halting it again.

- If you are not already grounded, properly ground yourself.
- Turn off the power supplies and disconnect the power, using the correct method for your system and power-supply type:

If your system uses...	Then...
AC power supplies	Unplug the power cords from the power source, and then remove the power cords.
DC power supplies	Remove the power at the DC source, and then remove the DC wires, if necessary.

Installing and cabling the new controller module

You must physically install the new controller module in the chassis, and then cable it.

Procedure

- If you have an I/O expansion module (IOXM) in your system and are creating a single-chassis HA pair, you must uncable and remove the IOXM.
You can then use the empty bay for the new controller module. However, the new configuration will not have the extra I/O provided by the IOXM.
- Physically install the new controller module and, if necessary, install additional fans:

If you are adding a controller module...	Then perform these steps...
To an empty bay to create a single-chassis HA pair and the system belongs to one of the following platforms:	<ol style="list-style-type: none"> a. Remove the blank plate in the rear of the chassis that covers the empty bay that will contain the new controller module. b. Gently push the controller module halfway into the chassis. <p>To prevent the controller module from automatically booting, do not fully seat it in the chassis until later in this procedure.</p>
In a separate chassis from its HA partner to create a dual-chassis HA pair when the existing configuration is in a controller-IOX module configuration. <ul style="list-style-type: none"> • HX6100 	Install the new system in the rack.

3. Cable the cluster network connections, as necessary:

- a) Identify the ports on the controller module for the cluster connections.

[AX2100/HX2000 Systems Installation and Setup Instructions](#)

[HX6100 Systems Installation and Setup Instructions](#)

- b) If you are configuring a switched cluster, identify the ports that you will use on the cluster network switches.

See the *Clustered Data ONTAP Switch Setup Guide for Cisco Switches*, *Fujitsu 10G Cluster-Mode Switch Installation Guide* or *Fujitsu 1G Cluster-Mode Switch Installation Guide*, depending on what switches you are using.
- c) Connect cables to the cluster ports:

If the cluster is...	Then...
A two-node switchless cluster	Directly connect the cluster ports on the existing controller module to the corresponding cluster ports on the new controller module.
A switched cluster	Connect the cluster ports on each controller to the ports on the cluster network switches identified in substep b.

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

The new controller module's FC-VI ports and HBAs (host bus adapters) must be cabled to the site FC switches.

Procedure

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

- **Port assignments for FC switches when using ONTAP 9.7 and later**
- Port assignments for systems using two initiator ports

Cabling the new controller module's cluster peering connections

You must cable the new controller module to the cluster peering network so that it has connectivity with the cluster on the partner site.

About this task

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

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

Procedure

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

Powering up both controller modules and displaying the LOADER prompt

You power up the existing controller module and the new controller module to display the LOADER prompt.

Procedure

Power up the controller modules and interrupt the boot process, using the steps for your configuration:

If the controller modules are...	Then...
In the same chassis	<ol style="list-style-type: none"> a. Verify that the new controller module is not fully inserted into the bay. The existing controller module should be fully inserted into the bay because it was never removed from the chassis, but the new controller module should not be. b. Connect the power and turn on the power supplies so that the existing controller module receives power. c. Interrupt the boot process on the existing controller module by pressing Ctrl-C. d. Push the new controller module firmly into the bay. When fully seated, the new controller module receives power and automatically boots. e. Interrupt the boot process by pressing Ctrl-C. f. Tighten the thumbscrew on the cam handle, if present. g. Install the cable management device, if present. h. Bind the cables to the cable management device with the hook and loop strap.
In separate chassis	<ol style="list-style-type: none"> a. Turn on the power supplies on the existing controller module. b. Interrupt the boot process by pressing Ctrl-C. c. Repeat these steps for the new controller module

Each controller module should display the LOADER prompt (LOADER>, LOADER-A>, or LOADER-B>).

Note: If there is no LOADER prompt, record the error message and contact Fujitsu support. If the system displays the boot menu, reboot and attempt to interrupt the boot process again.

Changing the ha-config setting on the existing and new controller modules

When you expand a MetroCluster configuration, you must update the ha-config setting of the existing controller module and the new controller module. You must also determine the system ID of the new controller module.

About this task

This task is performed in Maintenance mode on both the existing and new controller modules.

Procedure

1. Change the ha-config setting of the existing controller module:
 - a) Display the ha-config setting of the existing controller module and chassis: `ha-config show`
The ha-config setting is `mcc-2n` for all components because the controller module was in a two-node MetroCluster configuration.
 - b) Change the ha-config setting of the existing controller module to `mcc`: `ha-config modify controller mcc`
 - c) Change the ha-config setting of the existing chassis to `mcc`: `ha-config modify chassis mcc`
 - d) Retrieve the system ID for the existing controller module: `sysconfig`
Note the system ID. You need it when you set the partner ID on the new controller module.
 - e) Exit Maintenance mode to return to the LOADER prompt: `halt`
2. Change the ha-config setting and retrieve the system ID of the new controller module:
 - a) If the new controller module is not already in Maintenance mode, boot it to Maintenance mode: `boot_ontap maint`
 - b) Change the ha-config setting of the new controller module to `mcc`: `ha-config modify controller mcc`
 - c) Change the ha-config setting of the new chassis to `mcc`: `ha-config modify chassis mcc`
 - d) Retrieve the system ID for the new controller module: `sysconfig`
Note the system ID. You need it when you set the partner ID and assign disks to the new controller module.
 - e) Exit Maintenance mode to return to the LOADER prompt: `halt`

Setting the partner system ID for both controller modules

You must set the partner system ID on both controller modules so that they can form an HA pair.

About this task

This task is performed with both controller modules at the LOADER prompt.

Procedure

1. On the existing controller module, set the partner system ID to that of the new controller module: `setenv partner-sysid sysID_of_new_controller`
2. On the new controller module, set the partner system ID to that of the existing controller module: `setenv partner-sysid sysID_of_existing_controller`

Booting the existing controller module

You must boot the existing controller module to ONTAP.

Procedure

At the LOADER prompt, boot the existing controller module to ONTAP: `boot_ontap`

Assigning disks to the new controller module

Before you complete the configuration of the new controller module through netboot, you must assign disks to it.

Before you begin

You must have made sure that there are enough spares, unassigned disks, or assigned disks that are not part of an existing aggregate.

[Preparing for the upgrade](#) on page 85

About this task

These steps are performed on the existing controller module.

Procedure

1. Assign the root disk to the new controller module: `storage disk assign -disk disk_name -sysid new_controller_sysID -force true`
If your platform model uses the Advanced Drive Partitioning (ADP) feature, you must include the `-root true` parameter: `storage disk assign -disk disk_name -root true -sysid new_controller_sysID -force true`
2. Assign the remaining required disks to the new controller module by entering the following command for each disk: `storage disk assign -disk disk_name -sysid new_controller_sysID -force true`
3. Verify that the disk assignments are correct: `storage disk show -partitionownership`
Note: Ensure that you have assigned all disks that you intend to assign to the new node.

Netbooting and setting up ONTAP on the new controller module

You must perform a specific sequence of steps to netboot and install the ONTAP operating system on the new controller module when adding controller modules to an existing MetroCluster configuration.

About this task

- This task starts at the LOADER prompt of the new controller module.
- This task includes initializing disks.

The amount of time you need to initialize the disks depends on the size of the disks.

- The system automatically assigns two disks to the new controller module.

[Disks and Aggregates Power Guide](#)

Procedure

1. At the LOADER prompt, configure the IP address of the new controller module based on DHCP availability:

If DHCP is...	Then enter the following command...
Available	<code>ifconfig e0M -auto</code>
Not available	<code>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></code> <i>filer_addr</i> is the IP address of the storage system. <i>netmask</i> is the network mask of the storage system. <i>gateway</i> is the gateway for the storage system. <i>dns_addr</i> is the IP address of a name server on your network. <i>dns_domain</i> is the Domain Name System (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name. Note: Other parameters might be necessary for your interface. For details, use the <code>help ifconfig</code> command at the LOADER prompt.

2. At the LOADER prompt, netboot the new node:

```
netboot
http://web_server_ip/path_to_web-accessible_directory/
<ontap_version>_image.tgz
```

The *path_to_the_web-accessible_directory* is the location of the downloaded *<ontap_version>_image.tgz* file.

3. Select the *Install new software first* option from the displayed menu.

This menu option downloads and installs the new ONTAP image to the boot device.

- You should enter *y* when prompted with the message that this procedure is not supported for nondisruptive upgrade on an HA pair.
- You should enter *y* when warned that this process replaces the existing ONTAP software with new software.
- You should enter the path as follows when prompted for the URL of the *image.tgz* file:
`http://path_to_the_web-accessible_directory/image.tgz`

4. Enter *y* when prompted regarding nondisruptive upgrade or replacement of the software.

5. Enter the path to the *image.tgz* file when prompted for the URL of the package.

```
What is the URL for the package? http://path_to_web-accessible_directory/
image.tgz
```

6. Enter *n* to skip the backup recovery when prompted to restore the backup configuration.

```
*****
*           Restore Backup Configuration           *
* This procedure only applies to storage controllers that *
* are configured as an HA pair.                    *
*                                                    *
* Choose Yes to restore the "varfs" backup configuration *
* from the SSH server. Refer to the Boot Device Replacement *
* guide for more details.                          *
* Choose No to skip the backup recovery and return to the *
* boot menu.                                        *
*****
Do you want to restore the backup configuration
now? {y|n} n
```

7. Enter *y* when prompted to reboot now.

```
The node must be rebooted to start using the newly installed software. Do
you want to
reboot now? {y|n} y
```

8. If necessary, select the option to *Clean configuration and initialize all disks* after the node has booted.

Because you are configuring a new controller module and the new controller module's disks are empty, you can respond *y* when the system warns you that this will erase all disks.

Note: The amount of time needed to initialize disks depends on the size of your disks and configuration.

9. After the disks are initialized and the *Cluster Setup* wizard starts, set up the node:

- a) Enter the node management LIF information on the console.

10. Log in to the node, and enter the *cluster setup* and then enter *join* when prompted to join the cluster.

```
Do you want to create a new cluster or join an existing cluster? {create,
join}: join
```

11. Respond to the remaining prompts as appropriate for your site.

The *Software Setup Guide* for your version of ONTAP contains additional details.

12. If the system is in a two-node switchless cluster configuration, create the cluster interfaces on the existing node using the `network interface create` command to create cluster LIFs on the cluster ports. The following is an example command for creating a cluster LIF on one of the node's cluster ports. The `-auto` parameter configures the LIF to use a link-local IP address.

```
cluster_A::> network interface create -vserver Cluster -lif clus1 -role
cluster -home-node node_A_1 -home-port e1a -auto true
```

13. After setup is complete, verify that the node is healthy and eligible to participate in the cluster: `cluster show`

The following example shows a cluster after the second node (cluster1-02) has been joined to it:

```
cluster_A::> cluster show
Node                Health  Eligibility
-----
node_A_1            true   true
node_A_2            true   true
```

You can access the **Cluster Setup** wizard to change any of the values you entered for the admin storage virtual machine (SVM) or node SVM by using the `cluster setup` command.

14. Confirm that you have four ports configured as cluster interconnects: `network port show`
The following example shows output for two controller modules in cluster_A:

```
cluster_A::> network port show
Node  Port      IPspace      Broadcast Domain  Link  MTU  Speed (Mbps)
-----
node_A_1
  e0a      Cluster    Cluster      up      9000  auto/1000
  e0b      Cluster    Cluster      up      9000  auto/1000
  e0c      Default    Default      up      1500  auto/1000
  e0d      Default    Default      up      1500  auto/1000
  e0e      Default    Default      up      1500  auto/1000
  e0f      Default    Default      up      1500  auto/1000
  e0g      Default    Default      up      1500  auto/1000
node_A_2
  e0a      Cluster    Cluster      up      9000  auto/1000
  e0b      Cluster    Cluster      up      9000  auto/1000
  e0c      Default    Default      up      1500  auto/1000
  e0d      Default    Default      up      1500  auto/1000
  e0e      Default    Default      up      1500  auto/1000
  e0f      Default    Default      up      1500  auto/1000
  e0g      Default    Default      up      1500  auto/1000
14 entries were displayed.
```

Mirroring the root aggregate on the new controller

You must mirror the root aggregate to provide data protection when you are adding a controller to a MetroCluster configuration.

About this task

This task must be performed on the new controller module.

Procedure

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.

Configuring intercluster LIFs

You must create intercluster LIFs on ports used for communication between the MetroCluster partner clusters. You can use dedicated ports or ports that also have data traffic.

Configuring intercluster LIFs on dedicated ports

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

Procedure

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
	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	Cluster	cluster01-01:e0a, cluster01-01:e0b, cluster01-02:e0a, cluster01-02:e0b
cluster01	Default	cluster01-01:e0c, cluster01-01:e0d, cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f, cluster01-02:e0e, cluster01-02:e0f
	intercluster01	cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

5. Create intercluster LIFs on the system SVM and assign them to the failover group.

Option	Description
In ONTAP 9.7 and later:	<code>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</code>

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:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-intercluster</code>

For complete command syntax, see the man page.

```
cluster01::> network interface show -service-policy default-intercluster
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper  Address/Mask Node          Port         Home
-----
cluster01
cluster01_icl01  up/up      192.168.1.201/24 cluster01-01 e0e         true
cluster01_icl02  up/up      192.168.1.202/24 cluster01-02 e0f         true
```

7. Verify that the intercluster LIFs are redundant:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-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
Vserver      Logical      Home          Failover      Failover
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.

Procedure

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:

Option	Description
In ONTAP 9.7 and later:	<code>network interface create -vserver system_SVM -lif LIF_name -service-policy default-intercluster -home-node node -home-port port -address port_IP -netmask netmask</code>

For complete command syntax, see the man page.

The following example creates intercluster LIFs `cluster01_icl01` and `cluster01_icl02`:

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl01 -service-policy default-intercluster -home-node cluster01-01 -home-port e0c -address 192.168.1.201 -netmask 255.255.255.0
```

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl02 -service-policy default-intercluster -home-node cluster01-02 -home-port e0c -address 192.168.1.202 -netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-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	e0c	true
	cluster01_icl02	up/up	192.168.1.202/24	cluster01-02	e0c	true

4. Verify that the intercluster LIFs are redundant:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-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 `e0c` port will fail over to the `e0d` port.

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

Vserver	Logical Interface	Home Node:Port	Failover Policy	Failover Group
cluster01				
	cluster01_icl01	cluster01-01:e0c	local-only	192.168.1.201/24
			Failover Targets:	cluster01-01:e0c, cluster01-01:e0d
	cluster01_icl02	cluster01-02:e0c	local-only	192.168.1.201/24
			Failover Targets:	cluster01-02:e0c, cluster01-02:e0d

Creating a mirrored data aggregate on each node

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

Before you begin

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.

About this task

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.

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.

[Disks and Aggregates Power Guide](#)

Procedure

1. Display a list of available spares: `storage disk show -spare -owner node_name`
2. Create the aggregate by using the `storage aggregate create -mirror true` command.

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

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include

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

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

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

The following command creates a mirrored aggregate with 10 disks:

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

3. Verify the RAID group and drives of your new aggregate: `storage aggregate show-status -aggregate aggregate-name`

Installing licenses for the new controller module

You must add licenses for the new controller module for any ONTAP services that require standard (node-locked) licenses. For features with standard licenses, each node in the cluster must have its own key for the feature.

Procedure

Issue the following command to install each license key: `system license add -license-code license_key`

The `license_key` is 28 digits in length.

Repeat this step for each required standard (node-locked) license.

Creating unmirrored data aggregates

You can optionally create unmirrored data aggregates for data that does not require the redundant mirroring provided by MetroCluster configurations.

Before you begin

- You should know what drives or array LUNs will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can verify that the correct drive type is selected.

About this task



Attention:

In MetroCluster IP configurations, remote unmirrored aggregates are not accessible after a switchover

Note: The unmirrored aggregates must be local to the node owning them.

- Drives and array LUNs are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.
- Aggregate names should conform to the naming scheme you determined when you planned your MetroCluster configuration.
- The *Disks and Aggregates Power Guide* contains more information about mirroring aggregates.

Procedure

1. Enable unmirrored aggregate deployment: `metrocluster modify -enable-unmirrored-aggr-deployment true`

2. Verify that disk autoassignment is disabled: `disk option show`

3. Install and cable the disk shelves that will contain the unmirrored aggregates.

You can use the procedures in the Installation and Setup documentation for your platform and disk shelves.

[ETERNUS AX and ETERNUS HX Documentation Center](#)

4. Manually assign all disks on the new shelf to the appropriate node: `disk assign -disk disk-id -owner owner-node-name`

5. Create the aggregate: `storage aggregate create`

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

You must also ensure that you are only including drives on the unmirrored shelf to the aggregate.

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives or array LUNs that are to be added to the aggregate
- Number of drives to include
- Checksum style to use for the aggregate
- Type of drives to use
- Size of drives to use
- Drive speed to use
- RAID type for RAID groups on the aggregate
- Maximum number of drives or array LUNs that can be included in a RAID group

- Whether drives with different RPM are allowed

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

The following command creates a unmirrored aggregate with 10 disks:

```
controller_A_1::> storage aggregate create aggr1_controller_A_1 -
diskcount 10 -node controller_A_1
[Job 15] Job is queued: Create aggr1_controller_A_1.
[Job 15] The job is starting.
[Job 15] Job succeeded: DONE
```

6. Verify the RAID group and drives of your new aggregate: `storage aggregate show-status - aggregate aggregate-name`
7. Disable unmirrored aggregate deployment: `metrocluster modify -enable-unmirrored-aggr-deployment false`
8. Verify that disk autoassignment is enabled: `disk option show`

Related information

[Disks and Aggregates Power Guide](#)

Installing the firmware after adding a controller module

After adding the controller module, you must install the latest firmware on the new controller module so that the controller module functions properly with ONTAP.

Procedure

Download the most current version of firmware for your system and follow the instructions for downloading and installing the new firmware.

[Fujitsu Downloads: System Firmware and Diagnostics](#)

Refreshing the MetroCluster configuration with new controllers

You must refresh the MetroCluster configuration when expanding it from a two-node configuration to a four-node configuration.

Procedure

1. Refresh the MetroCluster configuration:
 - a) Enter advanced privilege mode: `set -privilege advanced`
 - b) Refresh the MetroCluster configuration: `metrocluster configure -refresh true`
The following command refreshes the MetroCluster configuration on all of the nodes in the DR group that contains controller_A_1:

```
controller_A_1::*> metrocluster configure -refresh true
```

```
[Job 726] Job succeeded: Configure is successful.
```

- c) Return to admin privilege mode: `set -privilege admin`
2. Verify the networking status on site A: `network port show`
The following example shows the network port usage on a four-node MetroCluster configuration:

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

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

```

e0g      Default  Default  up      1500  auto/1000
controller_A_2
e0a      Cluster  Cluster  up      9000  auto/1000
e0b      Cluster  Cluster  up      9000  auto/1000
e0c      Default  Default  up      1500  auto/1000
e0d      Default  Default  up      1500  auto/1000
e0e      Default  Default  up      1500  auto/1000
e0f      Default  Default  up      1500  auto/1000
e0g      Default  Default  up      1500  auto/1000

```

14 entries were displayed.

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

a) Verify the configuration from site A: `metrocluster show`

```

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

```

b) Verify the configuration from site B: `metrocluster show`

```

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

```

c) Verify that the DR relationships have been created correctly: `metrocluster node show -fields dr-cluster,dr-auxiliary,node-object-limit,automatic-uso,ha-partner,dr-partner`

```

metrocluster node show -fields dr-cluster,dr-auxiliary,node-object-limit,automatic-uso,ha-partner,dr-partner
dr-group-id cluster node ha-partner dr-cluster dr-partner dr-auxiliary node-object-limit automatic-uso
-----
2 cluster_A node_A_1 node_A_2 cluster_B node_B_1 node_B_2 on true
2 cluster_A node_A_2 node_A_1 cluster_B node_B_2 node_B_1 on true
2 cluster_B node_B_1 node_B_2 cluster_A node_A_1 node_A_2 on true
2 cluster_B node_B_2 node_B_1 cluster_A node_A_2 node_A_1 on true
4 entries were displayed.

```

Enabling storage failover on both controller modules and enabling cluster HA

After adding new controller modules to the MetroCluster configuration, you must enable storage failover on both controller modules and separately enable cluster HA.

Before you begin

The MetroCluster configuration must have previously been refreshed using the `metrocluster configure -refresh true` command.

About this task

This task must be performed on each MetroCluster site.

Procedure

1. Enable storage failover: `storage failover modify -enabled true -node existing-node-name`

The single command enables storage failover on both controller modules.

2. Verify that storage failover is enabled: `storage failover show`

The output should be similar to the following:

Node	Partner	Possible State	Description
old-ctlr	new-ctlr	true	Connected to new-ctlr
new-ctlr	old-ctlr	true	Connected to old-ctlr

2 entries were displayed.

3. Enable cluster HA: `cluster ha modify -configured true`

Cluster high availability (HA) must be configured in a cluster if it contains only two nodes and it differs from the HA provided by storage failover.

Restarting the SVMs

After expanding the MetroCluster configuration, you must restart the SVMs.

Procedure

1. Identify the SVMs that need to be restarted: `metrocluster vserver show`
This command shows the SVMs on both MetroCluster clusters.
2. Restart the SVMs on the first cluster:
 - a) Enter advanced privilege mode, pressing `y` when prompted: `set -privilege advanced`
 - b) Restart the SVMs: `vserver start -vserver SVM_name -force true`
 - c) Return to admin privilege mode: `set -privilege admin`
3. Repeat the previous step on the partner cluster.
4. Verify that the SVMs are in a healthy state: `metrocluster vserver show`

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.


Before you begin

- 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.
- 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.

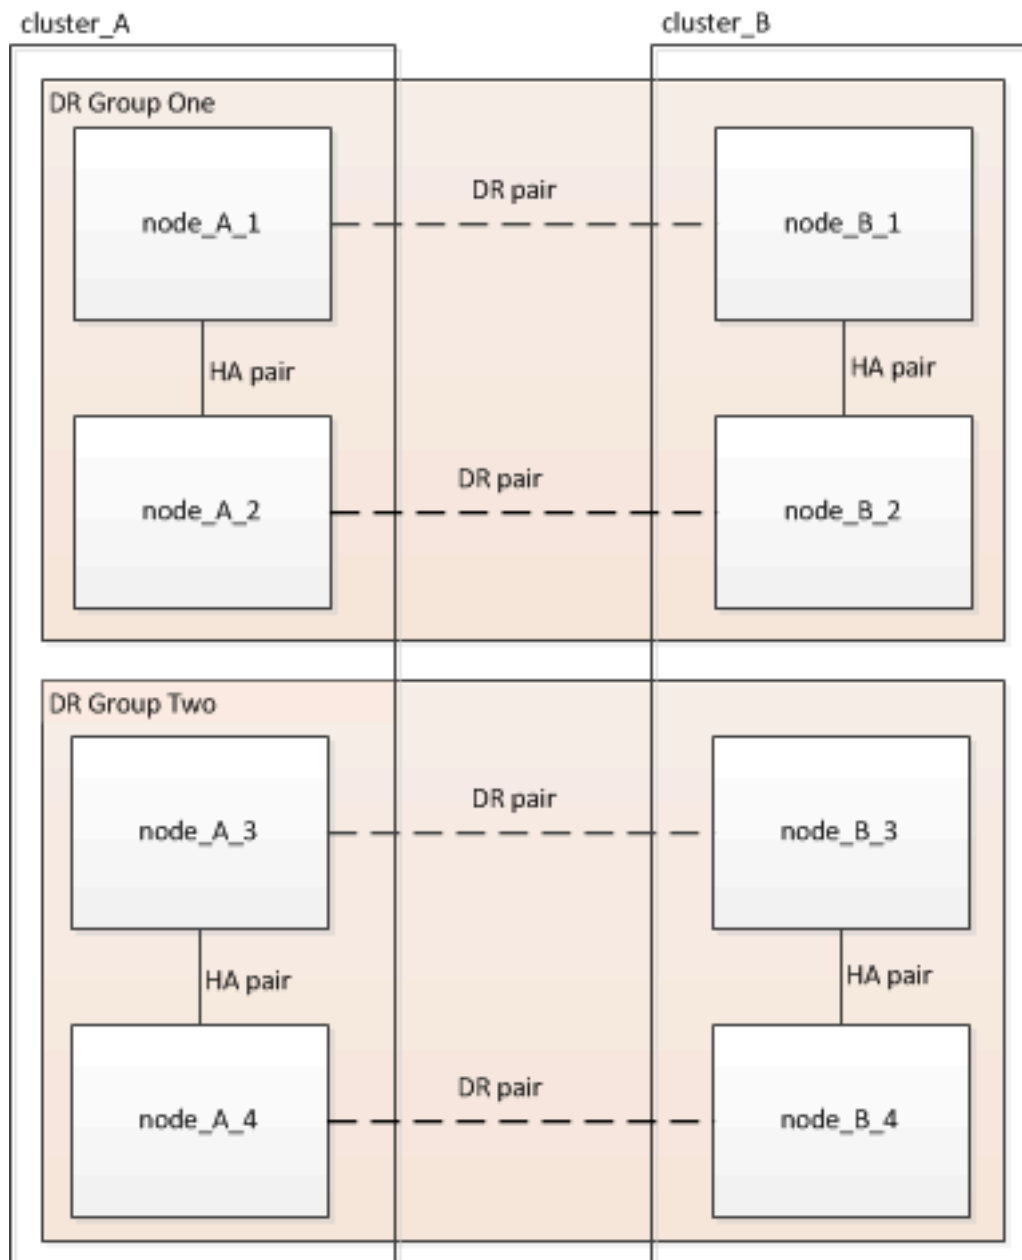
About this task

- 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.

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.

Procedure

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 guide so that the Reference Configuration Files (RCFs) can be used.

Note: Do not use this procedure if the cabling cannot use RCF files.

Racking the new equipment

You must rack the equipment for the new nodes.

Procedure

Use the MetroCluster Installation and Configuration guide and rack the new storage systems, disk shelves, and FC-to-SAS bridges.

Verifying the health of the MetroCluster configuration

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

Procedure

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

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

2. Check that mirroring is enabled on each node: `metrocluster node show`

```
cluster_A::> metrocluster node show
DR      Configuration  DR
Group Cluster Node      State      Mirroring Mode
-----
1      cluster_A
         node_A_1    configured  enabled   normal
         cluster_B
         node_B_1    configured  enabled   normal
2 entries were displayed.
```

3. Check that the MetroCluster components are healthy: `metrocluster check run`

```
cluster_A::> metrocluster check run

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

Component      Result
-----
nodes          ok
lifs           ok
config-replication ok
aggregates    ok
4 entries were displayed.
```

Command completed. Use the "metrocluster check show -instance" command or sub-commands in "metrocluster check" directory for detailed results. To check if the nodes are ready to do a switchover or switchback operation, run "metrocluster switchover -simulate" or "metrocluster switchback -simulate", respectively.

4. Check that there are no health alerts: `system health alert show`
5. Simulate a switchover operation:
 - a) From any node's prompt, change to the advanced privilege level: `set -privilege advanced`
You need to respond with `y` when prompted to continue into advanced mode and see the advanced mode prompt (`*>`).
 - b) Perform the switchover operation with the `-simulate` parameter: `metrocluster switchover -simulate`
 - c) Return to the admin privilege level: `set -privilege admin`

Checking for MetroCluster configuration errors with Config Advisor

You can go to the Fujitsu Support Site and download the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.

Note: Support for Config Advisor is limited, and available only online.

Procedure

1. Go to the Config Advisor download page and download the tool.
2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

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

You should issue an AutoSupport message to notify Fujitsu support that maintenance is underway. Informing Fujitsu 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.

Procedure

1. Log in to the cluster at Site_A.
2. Invoke an AutoSupport message indicating the start of the maintenance: `system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours`
maintenance-window-in-hours specifies the length of the maintenance window and can be a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can issue a command to indicating that the maintenance period has ended: `system node autosupport invoke -node * -type all -message MAINT=end`
3. Repeat this step on the partner site.

Recabling and zoning a switch fabric for the new nodes

When adding nodes to the MetroCluster configuration, you must change the cabling and then run RCF files to redefine the zoning on the fabric.

About this task

This task must be performed on each switch fabric. It is done one fabric at a time.

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.

Procedure

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 recable them. Only the ISL ports should remain cabled.

Applying the RCF files and recabling the switches

You must apply the RCF files to reconfigure your zoning to accommodate the new nodes.

Procedure

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.

Note: If your environment cannot be cabled in such a way that RCF files can be used then contact Fujitsu support. Do NOT use this procedure if the cabling cannot use RCF files.

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. Cable the FC-VI ports from the existing and new controllers, using the cabling layout you created previously.

Note: If your environment cannot be cabled in such a way that RCF files can be used then contact Fujitsu support. Do NOT use this procedure if the cabling cannot use RCF files.

8. From the existing nodes, verify that the FC-VI ports are online: `metrocluster interconnect adapter showmetrocluster 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 recable the second switch fabric.

Configuring ONTAP on the new controllers

You must set up ONTAP on each new controller in the MetroCluster configuration, and then re-create the MetroCluster relationship between the two sites.

Restoring system defaults on a previously used controller module

If your controller modules have been used previously, you must reset them for a successful MetroCluster configuration.

About this task

Important: This task is required only on controller modules that have been previously configured. You do not need to perform this task if you received the controller modules from the factory.

Procedure

1. At the LOADER prompt, return the environmental variables to their default setting: `set-defaults`
2. Boot the node to the boot menu: `boot_ontap menu`
After you run the command, wait until the boot menu is shown.
3. Clear the node configuration:
 - If you are using systems configured for ADP, select option 9a from the boot menu, and respond `yes` when prompted.

Note: This process is disruptive.

The following screen shows the boot menu prompt:

```
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 9a
##### WARNING #####

This is a disruptive operation and will result in the
loss of all filesystem data. Before proceeding further,
make sure that:
1) This option (9a) has been executed or will be executed
on the HA partner node, prior to reinitializing either
system in the HA-pair.
2) The HA partner node is currently in a halted state or
at the LOADER prompt.

Do you still want to continue (yes/no)? yes
```

- If your system is not configured for ADP, type `wipeconfig` at the boot menu prompt, and then press Enter.

The following screen shows the boot menu prompt:

```
Please choose one of the following:

(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? wipeconfig
This option deletes critical system configuration, including cluster
membership.
Warning: do not run this option on a HA node that has been taken
over.
```

```
Are you sure you want to continue?: yes
Rebooting to finish wipeconfig request.
```

Assigning disk ownership in ETERNUS AX systems

If you are using ETERNUS AX systems in a configuration with mirrored aggregates and the nodes do not have the disks (SSDs) correctly assigned, you should assign half the disks on each shelf to one local node and the other half of the disks to its HA partner node. You should create a configuration in which each node has the same number of disks in its local and remote disk pools.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

This does not apply to configurations which have unmirrored aggregates, an active/passive configuration, or that have an unequal number of disks in local and remote pools.

This task is not required if disks were correctly assigned when received from the factory.

Note: 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.

Procedure

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

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

- a) On the first node, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0: `disk assign -disk disk-name -p pool -n number-of-disks`
If storage controller Controller_A_1 has four shelves, each with 8 SSDs, you issue the following commands:

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

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

- b) Repeat the process for the second node at the local site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1: `disk assign -disk disk-name -p pool`
If storage controller Controller_A_1 has four shelves, each with 8 SSDs, you issue the following commands:

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

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

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

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

- a) On the first node at the remote site, systematically assign half the disks on each shelf to pool 0 and the other half to the HA partner's pool 0: `disk assign -disk disk-name -p pool`
If storage controller Controller_B_1 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0 -n 4

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1 -n 4
```

- b) Repeat the process for the second node at the remote site, systematically assigning half the disks on each shelf to pool 1 and the other half to the HA partner's pool 1: `disk assign -disk disk-name -p pool`
 If storage controller Controller_B_2 has four shelves, each with 8 SSDs, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0 -n 4
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0 -n 4

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1 -n 4
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1 -n 4
```

4. Confirm the disk assignments: `storage show disk`
5. Exit Maintenance mode: `halt`
6. Display the boot menu: `boot_ontap menu`
7. On each node, select option 4 to initialize all disks.

Assigning disk ownership in non-ETERNUS AX systems

If the MetroCluster nodes do not have the disks correctly assigned, or if you are using DS460C disk shelves in your configuration, you must assign disks to each of the nodes in the MetroCluster configuration on a shelf-by-shelf basis. You will create a configuration in which each node has the same number of disks in its local and remote disk pools.

Before you begin

The storage controllers must be in Maintenance mode.

About this task

If your configuration does not include DS460C disk shelves, this task is not required if disks were correctly assigned when received from the factory.

Note: Pool 0 always contains the disks that are found at the same site as the storage system that owns them.

Pool 1 always contains the disks that are remote to the storage system that owns them.

If your configuration includes DS460C disk shelves, you should manually assign the disks using the following guidelines for each 12-disk drawer:

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

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

Procedure

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-names shelf-name -p pool`

If storage controller Controller_B_1 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf1 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf2 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf1 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf2 -p 1
```

- b) Repeat the process for the second node at the remote site, systematically assigning its local disk shelves to pool 0 and its remote disk shelves to pool 1: `disk assign -shelf shelf-name -p pool`

If storage controller Controller_B_2 has four shelves, you issue the following commands:

```
*> disk assign -shelf FC_switch_B_1:1-5.shelf3 -p 0
*> disk assign -shelf FC_switch_B_1:1-5.shelf4 -p 0

*> disk assign -shelf FC_switch_A_1:1-5.shelf3 -p 1
*> disk assign -shelf FC_switch_A_1:1-5.shelf4 -p 1
```

4. Confirm the shelf assignments: `storage show shelf`

5. Exit Maintenance mode: `halt`

6. Display the boot menu: `boot_ontap menu`

7. On each node, select option 4 to initialize all disks.

Verifying the ha-config state of components

In a MetroCluster configuration, the ha-config state of the controller module and chassis components must be set to `mcc` so they boot up properly.

Before you begin

The system must be in Maintenance mode.

About this task

This task must be performed on each new controller module.

Procedure

1. In Maintenance mode, display the HA state of the controller module and chassis: `ha-config show`
The HA state for all components should be `mcc`.
2. If the displayed system state of the controller is not correct, set the HA state for the controller module: `ha-config modify controller mcc`

3. If the displayed system state of the chassis is not correct, set the HA state for the chassis: `ha-config modify chassis mcc`
4. Repeat these steps on the other replacement node.

Booting the new controllers and joining them to the cluster

To join the new controllers to the cluster, you must boot each new controller module and use the ONTAP cluster setup wizard to identify the cluster will join.

Before you begin

You must have cabled the MetroCluster configuration.

You must not have configured the Service Processor prior to performing this task.

About this task

This task must be performed on each of the new controllers at both clusters in the MetroCluster configuration.

Procedure

1. If you have not already done so, power up each node and let them boot completely.
If the system is in Maintenance mode, issue the `halt` command to exit Maintenance mode, and then issue the following command from the LOADER prompt: `boot_ontap`
The controller module enters the node setup wizard.
The output should be similar to the following:

```
Welcome to node setup

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the setup wizard.
  Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value.
.
.
.
```

2. Enable the AutoSupport tool by following the directions provided by the system.
3. Respond to the prompts to configure the node management interface.
The prompts are similar to the following:

```
Enter the node management interface port: [e0M]:
Enter the node management interface IP address: 10.228.160.229
Enter the node management interface netmask: 225.225.252.0
Enter the node management interface default gateway: 10.228.160.1
```

4. Confirm that nodes are configured in high-availability mode: `storage failover show -fields mode`

If not, you must issue the following command on each node, and then reboot the node: `storage failover modify -mode ha -node localhost`

This command configures high availability mode but does not enable storage failover. Storage failover is automatically enabled when you issue the `metrocluster configure` command later in the configuration process.

5. Confirm that you have four ports configured as cluster interconnects: `network port show`
The following example shows output for two controllers in cluster_A. If it is a two-node MetroCluster configuration, the output shows only one node.

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

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) 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.

Configuring the clusters into a MetroCluster configuration

You must peer the clusters, mirror the root aggregates, create a mirrored data aggregate, and then issue the command to implement the MetroCluster operations.

Configuring intercluster LIFs

You must create intercluster LIFs on ports used for communication between the MetroCluster partner clusters. You can use dedicated ports or ports that also have data traffic.

Configuring intercluster LIFs on dedicated ports

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

Procedure

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
	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

- 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

- Create a failover group for the dedicated ports:`network interface failover-groups create -vserver system_SVM -failover-group failover_group -targets physical_or_logical_ports`

The following example assigns ports e0e and e0f to the failover group intercluster01 on the system SVMcluster01:

```
cluster01::> network interface failover-groups create -vserver cluster01 -failover-group
intercluster01 -targets
cluster01-01:e0e,cluster01-01:e0f,cluster01-02:e0e,cluster01-02:e0f
```

- Verify that the failover group was created:`network interface failover-groups show`

For complete command syntax, see the man page.

```
cluster01::> network interface failover-groups show
```

Vserver	Group	Failover Targets

Cluster	Cluster	cluster01-01:e0a, cluster01-01:e0b, cluster01-02:e0a, cluster01-02:e0b
cluster01	Default	cluster01-01:e0c, cluster01-01:e0d, cluster01-02:e0c, cluster01-02:e0d, cluster01-01:e0e, cluster01-01:e0f, cluster01-02:e0e, cluster01-02:e0f
	intercluster01	cluster01-01:e0e, cluster01-01:e0f cluster01-02:e0e, cluster01-02:e0f

- Create intercluster LIFs on the system SVM and assign them to the failover group.

Option	Description
In ONTAP 9.7 and later:	
	<code>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</code>

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:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-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:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-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
Vserver      Logical Interface      Home Node:Port      Failover Policy      Failover 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.

Procedure

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      Cluster      up       1500     auto/1000
e0b   Cluster   Cluster      Cluster      up       1500     auto/1000
e0c   Default   Default      Default      up       1500     auto/1000
e0d   Default   Default      Default      up       1500     auto/1000
cluster01-02
e0a   Cluster   Cluster      Cluster      up       1500     auto/1000
e0b   Cluster   Cluster      Cluster      up       1500     auto/1000
e0c   Default   Default      Default      up       1500     auto/1000
e0d   Default   Default      Default      up       1500     auto/1000
```

2. Create intercluster LIFs on the system SVM:

Option	Description
In ONTAP 9.7 and later:	<code>network interface create -vserver system_SVM -lif LIF_name -service-policy default-intercluster -home-node node -home-port port -address port_IP -netmask netmask</code>

For complete command syntax, see the man page.

The following example creates intercluster LIFs `cluster01_icl01` and `cluster01_icl02`:

```
cluster01::> network interface create -vserver cluster01 -lif cluster01_icl01 -service-policy default-intercluster -home-node cluster01-01 -home-port e0c -address 192.168.1.201 -netmask 255.255.255.0

cluster01::> network interface create -vserver cluster01 -lif cluster01_icl02 -service-policy default-intercluster -home-node cluster01-02 -home-port e0c -address 192.168.1.202 -netmask 255.255.255.0
```

3. Verify that the intercluster LIFs were created:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-intercluster</code>

For complete command syntax, see the man page.

```
cluster01::> network interface show -service-policy default-intercluster
-----
Vserver      Logical      Status      Network      Current      Current      Is
Interface    Admin/Oper  Address/Mask Node          Port         Home
-----
cluster01
cluster01_icl01  up/up      192.168.1.201/24  cluster01-01  e0c         true
cluster01_icl02  up/up      192.168.1.202/24  cluster01-02  e0c         true
```

4. Verify that the intercluster LIFs are redundant:

Option	Description
In ONTAP 9.7 and later:	<code>network interface show -service-policy default-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 `e0c` port will fail over to the `e0d` port.

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

About this task

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

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

Note: 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.

Procedure

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.

Creating a mirrored data aggregate on each node

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

Before you begin

- You should know what drives will be used in the new aggregate.
- If you have multiple drive types in your system (heterogeneous storage), you should understand how you can ensure that the correct drive type is selected.

About this task

- Drives are owned by a specific node; when you create an aggregate, all drives in that aggregate must be owned by the same node, which becomes the home node for that aggregate.

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.

[Disks and Aggregates Power Guide](#)

Procedure

1. Display a list of available spares: `storage disk show -spare -owner node_name`
2. Create the aggregate by using the `storage aggregate create -mirror true` command.

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

You can specify the following options:

- Aggregate's home node (that is, the node that owns the aggregate in normal operation)
- List of specific drives that are to be added to the aggregate
- Number of drives to include

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

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

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

The following command creates a mirrored aggregate with 10 disks:

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

3. Verify the RAID group and drives of your new aggregate: `storage aggregate show-status - aggregate aggregate-name`

Implementing the MetroCluster configuration

You must run the `metrocluster configure -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.

Procedure

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) 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
-----
Local: cluster_B    Configuration state configured
                    Mode normal
Remote: cluster_A  Configuration state configured
                    Mode normal
```

Configuring FC-to-SAS bridges for health monitoring

About this task

- Third-party SNMP monitoring tools are not supported for FibreBridge bridges.
- Starting with ONTAP 9.8, FC-to-SAS bridges are monitored via in-band connections by default, and additional configuration is not required.

Note: Starting 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.

Procedure

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
9.7 and later	<code>storage bridge add -address 0.0.0.0 -managed-by in-band -name <i>bridge-name</i></code>

b) Verify that the bridge has been added and is properly configured: `storage bridge show`

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

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

```
controller_A_1::> storage bridge show

Bridge          Symbolic Name Is Monitored Monitor Status Vendor Model          Bridge WWN
-----
ATTO_10.10.20.10  att001        true         ok           Atto   FibreBridge 7500N
20000010867038c0
ATTO_10.10.20.11  att002        true         ok           Atto   FibreBridge 7500N
20000010867033c0
ATTO_10.10.20.12  att003        true         ok           Atto   FibreBridge 7500N
20000010867030c0
ATTO_10.10.20.13  att004        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.

Before you begin

- 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.

Procedure

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
```

Note: 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.

Procedure

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:

```
cluster_A::> metrocluster check aggregate show
```

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

Node	Aggregate	Check	Result
controller_A_1	controller_A_1_aggr0	mirroring-status	ok
		disk-pool-allocation	ok
		ownership-state	ok
	controller_A_1_aggr1	mirroring-status	ok
		disk-pool-allocation	ok
		ownership-state	ok
	controller_A_1_aggr2	mirroring-status	ok
		disk-pool-allocation	ok
		ownership-state	ok

```

controller_A_2      controller_A_2_aggr0      mirroring-status      ok
                    controller_A_2_aggr0      disk-pool-allocation  ok
                    controller_A_2_aggr0      ownership-state       ok
                    controller_A_2_aggr1      mirroring-status      ok
                    controller_A_2_aggr1      disk-pool-allocation  ok
                    controller_A_2_aggr1      ownership-state       ok
                    controller_A_2_aggr2      mirroring-status      ok
                    controller_A_2_aggr2      disk-pool-allocation  ok
                    controller_A_2_aggr2      ownership-state       ok
18 entries were displayed.

```

The following example shows the `metrocluster check cluster show` command output for a healthy four-node MetroCluster configuration. It indicates that the clusters are ready to perform a negotiated switchover if necessary.

```

Last Checked On: 9/13/2017 20:47:04
Cluster          Check                                     Result
-----
mccint-fas9000-0102
negotiated-switchover-ready      not-applicable
switchback-ready                 not-applicable
job-schedules                     ok
licenses                          ok
periodic-check-enabled           ok
mccint-fas9000-0304
negotiated-switchover-ready      not-applicable
switchback-ready                 not-applicable
job-schedules                     ok
licenses                          ok
periodic-check-enabled           ok
10 entries were displayed.

```

Checking for MetroCluster configuration errors with Config Advisor

You can go to the Fujitsu Support Site and download the Config Advisor tool to check for common configuration errors.

About this task

Config Advisor is a configuration validation and health check tool. You can deploy it at both secure sites and non-secure sites for data collection and system analysis.

Note: Support for Config Advisor is limited, and available only online.

Procedure

1. Go to the Config Advisor download page and download the tool.
2. Run Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

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

You should issue an AutoSupport message to notify Fujitsu support that maintenance is complete.

About this task

This task must be performed on each MetroCluster site.

Procedure

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.

Procedure

Use the procedures for negotiated switchover, healing, and switchback that are mentioned in the *MetroCluster Management and Disaster Recovery Guide*.

[MCC-MetroCluster Management and Disaster Recovery Guide](#)

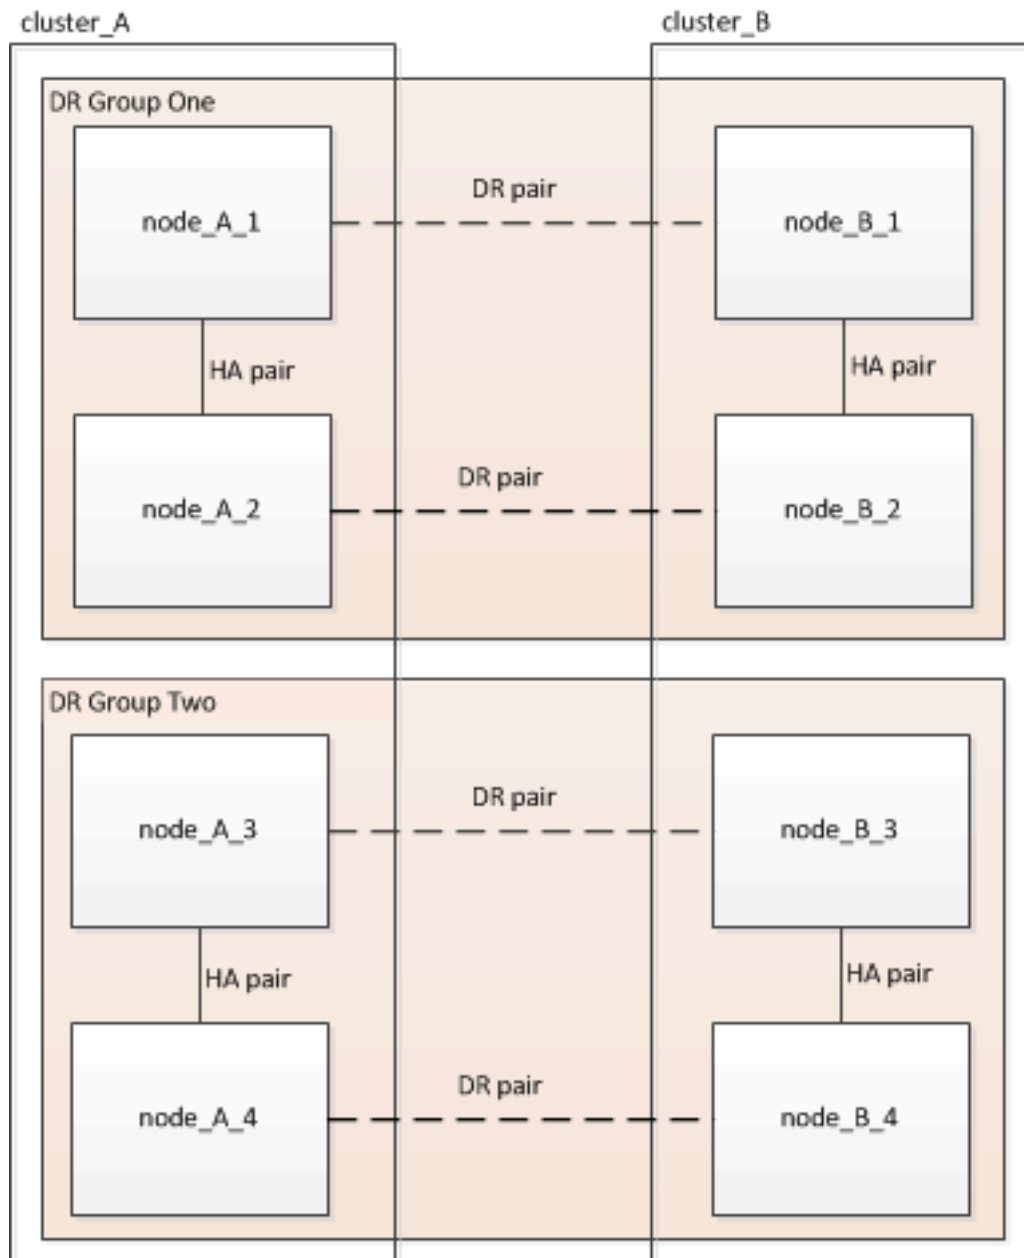
Removing a Disaster Recovery group

Starting 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. On earlier versions of ONTAP, please contact Fujitsu support to remove a DR group.

[Fujitsu Support](#)

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.



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.
- Prepare for the removal of the DR group, if you haven't already.
 - Move all data volumes to another DR group.
 - Move all MDV_CRS metadata volumes to another DR group. Follow the steps in the following procedure: [Moving a metadata volume in MetroCluster configurations](#)
 - Delete all MDV_aud metadata volumes that may exist in the DR group to be removed.
 - 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
```

Note: Root aggregates are not deleted.

- 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-port port-id`
- 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 data-lif -home-port port-id`

Node management and inter-cluster LIFs are not migrated.

- 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
```

- Identify the correct DR group for removal: `metrocluster node show`
 - 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. Unjoin the node: `cluster unjoin -node node-name`
- Repeat this step for the other local node in the old DR group.
- c. Set the admin privilege level: `set -privilege admin`
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.

Where to find additional information

You can learn more about MetroCluster configuration and operation from the Fujitsu documentation library.

MetroCluster and miscellaneous guides

Guide	Content
Fujitsu Manual Site	<ul style="list-style-type: none"> All MetroCluster guides
MetroCluster management and disaster recovery	<ul style="list-style-type: none"> Understanding the MetroCluster configuration Switchover, healing and switchback Disaster recovery
MetroCluster Service Guide	<ul style="list-style-type: none"> Guidelines for maintenance in a MetroCluster FC configuration Hardware replacement or upgrade and firmware upgrade procedures for FC-to-SAS bridges and FC switches Hot-adding a disk shelf in a fabric-attached or stretch MetroCluster FC configuration Hot-removing a disk shelf in a fabric-attached or stretch MetroCluster FC configuration Replacing hardware at a disaster site in a fabric-attached or stretch MetroCluster FC configuration Expanding a two-node fabric-attached or stretch MetroCluster FC configuration to a four-node MetroCluster configuration. Expanding a four-node fabric-attached or stretch MetroCluster FC configuration to an eight-node MetroCluster FC configuration.
MetroCluster Upgrade, Transition, and Expansion Guide	<ul style="list-style-type: none"> Upgrading or refreshing a MetroCluster configuration Transitioning from a MetroCluster FC configuration to a MetroCluster IP configuration Expanding a MetroCluster configuration by adding additional nodes
MetroCluster Tiebreaker Software Installation and Configuration Guide	<ul style="list-style-type: none"> Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
ERTERNUS AX and ETERNUS HX Documentation Center Note: The standard storage shelf maintenance procedures can be used with MetroCluster IP configurations.	<ul style="list-style-type: none"> Hot-adding a disk shelf Hot-removing a disk shelf
7-Mode Transition Tool Copy-Based Transition Guide	<ul style="list-style-type: none"> Transitioning data from 7-Mode storage systems to clustered storage systems
Concepts Guide	<ul style="list-style-type: none"> How mirrored aggregates work

Copyright and trademark

Copyright

Copyright 2021 FUJITSU LIMITED. All rights reserved.

No part of this document covered by copyright may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system - without prior written permission of the copyright owner.

Software derived from copyrighted Fujitsu material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY FUJITSU "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL FUJITSU BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Fujitsu reserves the right to change any products described herein at any time, and without notice. Fujitsu assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by Fujitsu. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of Fujitsu.

Trademark

FUJITSU, FUJITSU logo, and ETERNUS are trademarks of Fujitsu. All other trademarks are the property of their respective owners.

<https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html>

How to send comments about documentation and receive update notifications

The latest version of this document and the latest information related to this device are available at the following site.

<https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html>

If necessary, refer to the manuals for your model.

FUJITSU Storage ETERNUS AX/HX Series
MetroCluster® Upgrade, Transition, and Expansion Guide

A3CA08733-A285-01

Date of issuance: March 2021

Issuance responsibility: FUJITSU LIMITED

- The content of this manual is subject to change without notice.

- This manual was prepared with the utmost attention to detail.

However, Fujitsu shall assume no responsibility for any operational problems as the result of errors, omissions, or the use of information in this manual.

- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.

- The content of this manual may not be reproduced or distributed in part or in its entirety without prior permission from Fujitsu.