



Disruptively transition from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

ONTAP MetroCluster

NetApp
March 29, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-metrocluster/transition/task_disruptively_transition_from_a_two_node_mcc_fc_to_a_four_node_mcc_ip_configuration.html on March 29, 2024. Always check docs.netapp.com for the latest.

Table of Contents


- Disruptively transition from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later) 1
 - Disruptively transitioning from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later) 1
 - Example naming in this procedure 2
 - Preparing for disruptive FC-to-IP transition 2
 - Transitioning the MetroCluster FC nodes 11
 - Connecting the MetroCluster IP controller modules 16
 - Configuring the new nodes and completing transition 32
 - Returning the system to normal operation 37

Disruptively transition from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Disruptively transitioning from a two-node MetroCluster FC to a four-node MetroCluster IP configuration (ONTAP 9.8 and later)

Beginning with ONTAP 9.8, you can transition workloads and data from an existing two-node MetroCluster FC configuration to a new four-node MetroCluster IP configuration. Disk shelves from the MetroCluster FC nodes are moved to the IP nodes.

The following illustration provides a simplified view of the configuration before and after this transition procedure.



- This procedure is supported on systems running ONTAP 9.8 and later.
- This procedure is disruptive.
- This procedure applies only to a two-node MetroCluster FC configuration.

If you have a four-node MetroCluster FC configuration, see [Choosing your transition procedure](#).

- ADP is not supported on the four-node MetroCluster IP configuration created by this procedure.
- You must meet all requirements and follow all steps in the procedure.
- The existing storage shelves are moved to the new MetroCluster IP nodes.
- Additional storage shelves can be added to the configuration if necessary.

Example naming in this procedure

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

The nodes in the original configuration have the suffix -FC, indicating that they are in a fabric-attached or stretch MetroCluster configuration.

Components	cluster_A at site_A	cluster_B at site_B
dr_group_1-FC	<ul style="list-style-type: none">• node_A_1-FC• shelf_A_1• shelf_A_2	<ul style="list-style-type: none">• node_B_1-FC• shelf_B_1• shelf_B_2
dr_group_2-IP	<ul style="list-style-type: none">• node_A_1-IP• node_A_2-IP• shelf_A_1• shelf_A_2• shelf_A_3-new• shelf_A_4-new	<ul style="list-style-type: none">• node_B_1-IP• node_B_2-IP• shelf_B_1• shelf_B_2• shelf_B_3-new• shelf_B_4-new
Switches	<ul style="list-style-type: none">• switch_A_1-FC• switch_A_2-FC• switch_A_1-IP• switch_A_2-IP	<ul style="list-style-type: none">• switch_B_1-FC• switch_B_2-FC• switch_B_1-IP• switch_B_2-IP

Preparing for disruptive FC-to-IP transition

General requirements for disruptive FC-to-IP transition

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

The existing MetroCluster FC configuration must meet the following requirements:

- It must be a two-node configuration and all nodes must be running ONTAP 9.8 or later.

It can be a two-node fabric-attached or stretched MetroCluster.

- It must meet all requirements and cabling as described in the *MetroCluster Installation and Configuration* procedures.

[Fabric-attached MetroCluster installation and configuration](#)

[Stretch MetroCluster installation and configuration](#)

- It cannot be configured with NetApp Storage Encryption (NSE).
- The MDV volumes cannot be encrypted.

You must have remote console access for all six nodes from either MetroCluster site or plan for travel between the sites as required by the procedure.

Drive shelf reuse and drive requirements for disruptive FC-to-IP transition

You must ensure that adequate spare drives and root aggregate space is available on the storage shelves.

Reusing the existing storage shelves

When using this procedure, the existing storage shelves are retained for use by the new configuration. When node_A_1-FC and node_B_1-FC are removed, the existing drive shelves are connected to node_A_1-IP and node_A_2-IP on cluster_A and to node_B_1-IP and node_B_2-IP on cluster_B.

- The existing storage shelves (those attached to node_A_1-FC and node_B_1-FC) must be supported by the new platform models.

If the existing shelves are not supported by the new platform models, see [Disruptively transitioning when existing shelves are not supported on new controllers \(ONTAP 9.8 and later\)](#).

- You must ensure you don't exceed the platform limits for drives, etc.


[NetApp Hardware Universe](#)

Storage requirements for the additional controllers

Additional storage must be added, if necessary, to accommodate the two additional controllers (node_A_2-IP and node_B_2-ip), because the configuration is changing from a two-node to a four-node arrangement.

- Depending on the spare drives available in the existing shelves, additional drives must be added to accommodate the additional controllers in the configuration.

This might require additional storage shelves, as shown in the following illustration.



You need to have additional 14 - 18 drives each for the third and fourth controllers (node_A_2-IP and node_B_2-IP):

- Three pool0 drives
- Three pool1 drives
- Two spare drives
- Six to ten drives for the system volume
- You must ensure that the configuration, including the new nodes, does not exceed the platform limits for the configuration, including drive count, root aggregate size capacity, etc.


This information is available for each platform model at *NetApp Hardware Universe*.

[NetApp Hardware Universe](#)

Workflow for disruptive transition

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

As you prepare for the transition, plan for travel between the sites. Note that after the remote nodes are racked and cabled, you need serial terminal access to the nodes. Service Processor access is not be available until the nodes are configured.



Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes

You must adjust the port and LIF configuration of the MetroCluster FC node so it is compatible with that of the MetroCluster IP node that will replace it.

About this task

When the new nodes are first booted during the upgrade process, each node uses the most recent configuration of the node it is replacing. When you boot node_A_1-IP, ONTAP attempts to host LIFs on the same ports that were used on node_A_1-FC.

During the transition procedure, you will perform steps on both the old and new nodes to ensure correct cluster, management, and data LIF configuration.

Steps

1. Identify any conflicts between the existing MetroCluster FC port usage and the port usage for the MetroCluster IP interfaces on the new nodes.

You must identify the MetroCluster IP ports on the new MetroCluster IP controllers using the table below. Then check and record if any data LIFs or cluster LIFs exist on those ports on the MetroCluster FC nodes.

These conflicting data LIFs or cluster LIFs on the MetroCluster FC nodes will be moved at the appropriate step in the transition procedure.

The following table shows the MetroCluster IP ports by platform model. You can ignore the VLAN ID column.

Platform model	MetroCluster IP port	VLAN ID	
----------------	----------------------	---------	--

AFF A800	e0b	Not used	
	e1b		
AFF A700 and FAS9000	e5a		
	e5b		
AFF A320	e0g		
	e0h		
AFF A300 and FAS8200	e1a		
	e1b		
FAS8300/A400/FAS8700	e1a	10	
	e1b	20	
AFF A250 and FAS500f	e0c	10	
	e0b	20	

You can fill in the following table and refer to it later in the transition procedure.

Ports	Corresponding MetroCluster IP interface ports (from table above)	Conflicting LIFs on these ports on the MetroCluster FC nodes
First MetroCluster IP port on node_A_1-FC		
Second MetroCluster IP port on node_A_1-FC		
First MetroCluster IP port on node_B_1-FC		
Second MetroCluster IP port on node_B_1-FC		

- Determine which physical ports are available on the new controllers and which LIFs can be hosted on the ports.

The controller's port usage depends on the platform model and IP switch model you will use in the MetroCluster IP configuration. You can gather the port usage of the new platforms from the *NetApp*

NetApp Hardware Universe

3. If desired, record the port information for node_A_1-FC and node_A_1-IP.

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

In the columns for node_A_1-IP, add the physical ports for the new controller module and plan the IPspaces and broadcast domains for the new node.

	node_A_1-FC			node_A_1-IP		
LIF	Ports	IPspaces	Broadcast domains	Ports	IPspaces	Broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

4. If desired, record all the port information for node_B_1-FC.

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

In the columns for node_B_1-IP, add the physical ports for the new controller module and plan the LIF port usage, IPspaces and broadcast domains for the new node.

	node_B_1-FC			node_B_1-IP		
LIF	Physical ports	IPspaces	Broadcast domains	Physical ports	IPspaces	Broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

Preparing the MetroCluster IP controllers

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

About this task

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

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

The nodes should be connected to any **new** storage shelves. They must **not** be connected to the existing storage shelves containing data.

These steps can be performed now, or later in the procedure when the controllers and shelves are racked. In any case, you must make sure you clear the configuration and prepare the nodes **before** connecting them to the existing storage shelves and **before** making any configuration changes to the MetroCluster FC nodes.



Do not perform these steps with the MetroCluster IP controllers connected to the existing storage shelves that were connected to the MetroCluster FC controllers.

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

Steps

1. Connect the controller modules to the new storage shelves.
2. In Maintenance mode, display the HA state of the controller module and chassis:

```
ha-config show
```

The HA state for all components should be “mccip”.

3. If the displayed system state of the controller or chassis is not correct, set the HA state:

```
ha-config modify controller mccip`ha-config modify chassis mccip
```

4. Exit Maintenance mode:

```
halt
```

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

5. Repeat the following substeps on all four nodes to clear the configuration:

- a. Set the environmental variables to default values:

```
set-defaults
```

- b. Save the environment:

```
saveenv
```

```
bye
```

6. Repeat the following substeps to boot all four nodes using the 9a option on the boot menu.

- a. At the LOADER prompt, launch the boot menu:

```
boot_ontap menu
```

- b. At the boot menu, select option “9a” to reboot the controller.

7. Boot each of the four nodes to Maintenance mode using option “5” on the boot menu.

8. Record the system ID and from each of the four nodes:

```
sysconfig
```

9. Repeat the following steps on node_A_1-IP and node_B_1-IP.

- a. Assign ownership of all disks local to each site:

```
disk assign adapter.xx.*
```

b. Repeat the previous step for each HBA with attached drive shelves on node_A_1-IP and node_B_1-IP.

10. Repeat the following steps on node_A_1-IP and node_B_1-IP to clear the mailbox region on each local disk.

a. Destroy the mailbox region on each disk:

```
mailbox destroy local ``mailbox destroy partner
```

11. Halt all four controllers:

```
halt
```

12. On each controller, display the boot menu:

```
boot_ontap menu
```

13. On each of the four controllers, clear the configuration:

```
wipeconfig
```

When the wipeconfig operation completes, the node automatically returns to the boot menu.

14. Repeat the following substeps to again boot all four nodes using the 9a option on the boot menu.

a. At the LOADER prompt, launch the boot menu:

```
boot_ontap menu
```

b. At the boot menu, select option “9a” to reboot the controller.

c. Let the controller module complete booting before moving to the next controller module.

After “9a” completes, the nodes automatically return to the boot menu.

15. Power off the controllers.

Verifying the health of the MetroCluster FC configuration

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

This task is performed on the MetroCluster FC configuration.

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

a. Check whether the system is multipathed:

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

b. Check for any health alerts on both clusters:

```
system health alert show
```

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

```
metrocluster show
```

- d. Perform a MetroCluster check:

```
metrocluster check run
```

- e. Display the results of the MetroCluster check:

```
metrocluster check show
```

- f. Check for any health alerts on the switches (if present):

```
storage switch show
```

- g. Run Config Advisor.

[NetApp Downloads: Config Advisor](#)

- h. After running Config Advisor, review the tool's output and follow the recommendations in the output to address any issues discovered.

2. Verify that the nodes are in non-HA mode:

```
storage failover show
```

Removing the existing configuration from the Tiebreaker or other monitoring software

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

Steps

1. Remove the existing MetroCluster configuration from the Tiebreaker software.

[Removing MetroCluster configurations](#)

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

Refer to the documentation for the application.

Transitioning the MetroCluster FC nodes

You must gather information from the existing MetroCluster FC nodes, send an autosupport message announcing the start of maintenance, and transition the nodes.

Gathering information from the existing controller modules before the transition

Before transitioning, you must gather information for each of the nodes.

This task is performed on the existing nodes:

- node_A_1-FC
- node_B_1-FC

1. Gather the output for the commands in the following table.

Category	Commands	Notes
License	system license show	
Shelves and numbers of disks in each shelf and flash storage details and memory and NVRAM and network cards	system node run -node node_name sysconfig	
Cluster network and node management LIFs	system node run -node node_name sysconfig network interface show -role "cluster,node-mgmt,data"	
SVM information	vserver show	
Protocol information	nfs show iscsi show cifs show	
Physical ports	network port show -node node_name -type physical network port show	
Failover Groups	network interface failover-groups show -vserver vserver_name	Record the names and ports of failover groups that are not clusterwide.
VLAN configuration	network port vlan show -node node_name	Record each network port and VLAN ID pairing.
Interface group configuration	network port ifgrp show -node node_name -instance	Record the names of the interface groups and the ports assigned to them.
Broadcast domains	network port broadcast-domain show	
IPspace	network ipspace show	
Volume info	volume show and volume show -fields encrypt	
Aggregate Info	storage aggregate show and storage aggr encryption show and storage aggregate object-store show	
Disk ownership information	storage aggregate show and storage aggr encryption show and storage aggregate object-store show	
Encryption	storage failover mailbox-disk show and security key-manager backup show	Also preserve the passphrase used to enable key-manager. In the case of external key-manager you will need the authentication information for the client and server.

Category	Commands	Notes
Encryption	security key-manager show	
Encryption	security key-manager external show	
Encryption	systemshell local kenv kmip.init.ipaddr ip-address	
Encryption	systemshell local kenv kmip.init.netmask netmask	
Encryption	systemshell local kenv kmip.init.gateway gateway	
Encryption	systemshell local kenv kmip.init.interface interface	

Sending a custom AutoSupport message prior to maintenance

Before performing the maintenance, you should issue an AutoSupport message to notify NetApp technical support that maintenance is underway. This prevents them from opening a case on the assumption that a disruption has occurred.

This task must be performed on each MetroCluster site.

1. To prevent automatic support case generation, send an Autosupport message to indicate maintenance is underway.

- a. Issue the following command: `system node autosupport invoke -node * -type all -message MAINT=maintenance-window-in-hours`

`maintenance-window-in-hours` specifies the length of the maintenance window, with a maximum of 72 hours. If the maintenance is completed before the time has elapsed, you can invoke an AutoSupport message indicating the end of the maintenance period: `system node autosupport invoke -node * -type all -message MAINT=end`

- b. Repeat the command on the partner cluster.

Transitioning, shutting down, and removing the MetroCluster FC nodes

In addition to issuing commands on the MetroCluster FC nodes, this task includes physical uncabling and removal of the controller modules at each site.

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

- node_A_1-FC
- node_B_1-FC
 1. Stop all client traffic.
 2. On either of the MetroCluster FC nodes, for example node_A_1-FC, enable transition.
 - a. Set the advanced privilege level: `set -priv advanced`

b. Enable transition: `metrocluster transition enable -transition-mode disruptive`

c. Return to admin mode: `set -priv admin`

3. Unmirror the root aggregate by deleting the remote plex of the root aggregates.

a. Identify the root aggregates: `storage aggregate show -root true`

b. Display the pool1 aggregates: `storage aggregate plex show -pool 1`

c. Delete the local plex of the root aggregate: `aggr plex delete aggr-name -plex plex-name`

d. Offline the remote plex of the root aggregate: `aggr plex offline root-aggregate -plex remote-plex-for-root-aggregate`

For example:

```
# aggr plex offline aggr0_node_A_1-FC_01 -plex plex4
```

4. Confirm the mailbox count, disk autoassign, and transition mode before proceeding using the following commands on each controller:

a. Set the advanced privilege level: `set -priv advanced`

b. Confirm that only three mailbox drives are shown for each controller module: `storage failover mailbox-disk show`

c. Return to admin mode: `set -priv admin`

d. Confirm that the transition mode is disruptive: `metrocluster transition show`

5. Check for any broken disks: `disk show -broken`

6. Remove or replace any broken disks

7. Confirm aggregates are healthy using the following commands on node_A_1-FC and node_B_1-FC: `storage aggregate show` `storage aggregate plex show`

The storage aggregate show command indicates that the root aggregate is unmirrored.

8. Check for any VLANs or interface groups: `network port ifgrp show` `network port vlan show`

If none are present, skip the following two steps.

9. Display the list of Lifs using VLANs or ifgrps: `network interface show -fields home-port,curr-port` `network port show -type if-group | vlan`

10. Remove any VLANs and interface groups.

You must perform these steps for all LIFs in all SVMs, including those SVMs with the -mc suffix.

a. Move any LIFs using the VLANs or interface groups to an available port: `network interface modify -vserver vservers-name -lif lif_name -home- port port`

b. Display the LIFs that are not on their home ports: `network interface show -is-home false`

- c. Revert all LIFs to their respective home ports: `network interface revert -vserver vserver_name -lif lif_name`
- d. Verify that all LIFs are on their home ports: `network interface show -is-home false`

No LIFs should appear in the output.

- e. Remove VLAN and ifgrp ports from broadcast domain: `network port broadcast-domain remove-ports -ip-space ip-space -broadcast-domain broadcast-domain-name -ports nodename:portname,nodename:portname,...`
 - f. Verify that all the vlan and ifgrp ports are not assigned to a broadcast domain: `network port show -type if-group | vlan`
 - g. Delete all VLANs: `network port vlan delete -node nodename -vlan-name vlan-name`
 - h. Delete interface groups: `network port ifgrp delete -node nodename -ifgrp ifgrp-name`
11. Move any LIFs as required to resolve conflicts with the MetroCluster IP interface ports.


You must move the LIFs identified in step 1 of [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#).

- a. Move any LIFs hosted on the desired port to another port: `network interface modify -lif lifname -vserver vserver-name -home-port new-homeport`network interface revert -lif lifname -vserver vservername`
 - b. If necessary, move the destination port to an appropriate IPspace and broadcast domain. `network port broadcast-domain remove-ports -ip-space current-ip-space -broadcast-domain current-broadcast-domain -ports controller-name:current-port`network port broadcast-domain add-ports -ip-space new-ip-space -broadcast-domain new-broadcast-domain -ports controller-name:new-port`
12. Halt the MetroCluster FC controllers (node_A_1-FC and node_B_1-FC): `system node halt`
13. At the LOADER prompt, synchronize the hardware clocks between the FC and IP controller modules.
- a. On the old MetroCluster FC node (node_A_1-FC), display the date: `show date`
 - b. On the new MetroCluster IP controllers (node_A_1-IP and node_B_1-IP), set the date shown on original controller: `set date mm/dd/yy`
 - c. On the new MetroCluster IP controllers (node_A_1-IP and node_B_1-IP), verify the date: `show date`
14. Halt and power off the MetroCluster FC controller modules (node_A_1-FC and node_B_1-FC), FC-to-SAS bridges (if present), FC switches (if present) and each storage shelf connected to these nodes.
15. Disconnect the shelves from the MetroCluster FC controllers and document which shelves are local storage to each cluster.

If the configuration uses FC-to-SAS bridges or FC back-end switches, disconnect and remove them.

16. In Maintenance mode on the MetroCluster FC nodes (node_A_1-FC and node_B_1-FC), confirm no disks are connected: `disk show -v`
17. Power down and remove the MetroCluster FC nodes.

At this point, the MetroCluster FC controllers have been removed and the shelves are disconnected from all controllers.



Connecting the MetroCluster IP controller modules

You must add the four new controller modules and any additional storage shelves to the configuration. The new controller modules are added two-at-a-time.

Setting up the new controllers


You must rack and cable the new MetroCluster IP controllers to the storage shelves previously connected to the MetroCluster FC controllers.

About this task

These steps must be performed on each of the MetroCluster IP nodes.

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

In the following example, two additional storage shelves are added at each site to provide storage to accommodate the new controller modules.



Steps

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

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

2. Properly ground yourself.
3. Rack the new equipment: controllers, storage shelves, and IP switches.

Do not cable the storage shelves or IP switches at this time.

4. Connect the power cables and management console connection to the controllers.
5. Verify that all storage shelves are powered off.
6. Verify that no drives are connected by performing the following steps on all four nodes:
 - a. At the LOADER prompt, launch the boot menu:

```
boot_ontap maint
```

- b. Verify that no drives are connected:

```
disk show -v
```

The output should show no drives.

- c. Halt the node:

```
halt
```

7. Boot all four nodes using the 9a option on the boot menu.

- a. At the LOADER prompt, launch the boot menu:

```
boot_ontap menu
```

- b. At the boot menu, select option “9a” to reboot the controller.
- c. Let the controller module complete booting before moving to the next controller module.

After “9a” completes, the nodes automatically return to the boot menu.

8. Cable the storage shelves.

Refer to the controller installation and setup procedures for your model for cabling information.

[ONTAP Hardware Systems Documentation](#)

9. Cable the controllers to the IP switches as described in [Cabling the IP switches](#).
10. Prepare the IP switches for the application of the new RCF files.

Follow the steps for your switch vendor:

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

11. Download and install the RCF files.

Follow the steps for your switch vendor:

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

12. Turn on power to the first new controller (node_A_1-IP) and press Ctrl-C to interrupt the boot process and display the LOADER prompt.
13. Boot the controller to Maintenance mode:

```
boot_ontap_maint
```

14. Display the system ID for the controller:

```
sysconfig -v
```

15. Confirm that the shelves from the existing configuration are visible from the new MetroCluster IP node:

```
storage show shelf``disk show -v
```

16. Halt the node:

```
halt
```

17. Repeat the preceding steps on the other node at the partner site (site_B).

Connecting and booting up node_A_1-IP and node_B_1-IP

After connecting the MetroCluster IP controllers and IP switches, you transition and boot up node_A_1-IP and node_B_1-IP.

Bringing up node_A_1-IP

You must boot the node with the correct transition option.

Steps

1. Boot node_A_1-IP to the boot menu:

```
boot_ontap menu
```

2. Issue the following command at the boot menu prompt to initiate transition:

```
boot_after_mcc_transition
```

- This command reassigns all the disks owned by node_A_1-FC to node_A_1-IP.
 - node_A_1-FC disks are assigned to node_A_1-IP
 - node_B_1-FC disks are assigned to node_B_1-IP
- The command also automatically makes other required system ID reassignments so the MetroCluster IP nodes can boot to the ONTAP prompt.
- If the boot_after_mcc_transition command fails for any reason, it should be re-run from the boot menu.



- If the following prompt is displayed, enter Ctrl-C to continue. Checking MCC DR state... [enter Ctrl-C(resume), S(status), L(link)]_
- If the root volume was encrypted, the node halts with the following message. Halting the system, because root volume is encrypted (NetApp Volume Encryption) and the key import failed. If this cluster is configured with external (KMIP) key-manager, check the health of the key servers.

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning. Selection (1-9)?

``boot_after_mcc_transition``

This will replace all flash-based configuration with the last backup to disks. Are you sure you want to continue?: yes

MetroCluster Transition: Name of the MetroCluster FC node: ``node_A_1-FC``

MetroCluster Transition: Please confirm if this is the correct value
[yes|no]:? y

MetroCluster Transition: Disaster Recovery partner sysid of
MetroCluster FC node node_A_1-FC: ``systemID-of-node_B_1-FC``

MetroCluster Transition: Please confirm if this is the correct value
[yes|no]:? y

MetroCluster Transition: Disaster Recovery partner sysid of local
MetroCluster IP node: ``systemID-of-node_B_1-IP``

MetroCluster Transition: Please confirm if this is the correct value
[yes|no]:? y

3. If data volumes are encrypted, restore the keys using the correct command for your key management configuration.

If you are using...	Use this command...
Onboard key management	<code>security key-manager onboard sync</code> For more information, see Restoring onboard key management encryption keys .
External key management	<code>security key-manager key query -node node-name</code> For more information, see Restoring external key management encryption keys .

4. If the root volume is encrypted, use the procedure in [Recovering key management if the root volume is encrypted](#).

Recovering key management if the root volume is encrypted

If the root volume is encrypted, you must use special boot commands to restore the key management.

Before you begin

You must have the passphrases gathered earlier.

Steps

1. If onboard key management is used, perform the following substeps to restore the configuration.
 - a. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- b. Select option “(10) Set onboard key management recovery secrets” from the boot menu.

Respond as appropriate to the prompts:

```
This option must be used only in disaster recovery procedures. Are
you sure? (y or n): y
Enter the passphrase for onboard key management: passphrase
Enter the passphrase again to confirm: passphrase

Enter the backup data: backup-key
```

The system boots to the boot menu.

- c. Enter option “6” at the boot menu.

Respond as appropriate to the prompts:

```
This will replace all flash-based configuration with the last backup
to
disks. Are you sure you want to continue?: y

Following this, the system will reboot a few times and the following
prompt will be available continue by saying y

WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
```

After the reboots, the system will be at the LOADER prompt.

- d. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- e. Again elect option “(10) Set onboard key management recovery secrets” from the boot menu.

Respond as appropriate to the prompts:

```
This option must be used only in disaster recovery procedures. Are
you sure? (y or n): `y`
Enter the passphrase for onboard key management: `passphrase`
Enter the passphrase again to confirm: `passphrase`

Enter the backup data: `backup-key`
```

The system boots to the boot menu.

- f. Enter option “1” at the boot menu.

If the following prompt is displayed, you can press Ctrl+C to resume the process.

```
Checking MCC DR state... [enter Ctrl-C(resume), S(status), L(link)]
```

The system boots to the ONTAP prompt.

- g. Restore the onboard key management:

```
security key-manager onboard sync
```

Respond as appropriate to the prompts, using the passphrase you collected earlier:

```
cluster_A::> security key-manager onboard sync
Enter the cluster-wide passphrase for onboard key management in
Vserver "cluster_A":: passphrase
```

- 2. If external key management is used, perform the following substeps to restore the configuration.

- a. Set the required bootargs:

```
setenv bootarg.kmip.init.ipaddr ip-address

setenv bootarg.kmip.init.netmask netmask

setenv bootarg.kmip.init.gateway gateway-address

setenv bootarg.kmip.init.interface interface-id
```

- b. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- c. Select option “(11) Configure node for external key management” from the boot menu.

The system boots to the boot menu.

- d. Enter option “6” at the boot menu.

The system boots multiple times. You can respond affirmatively when prompted to continue the boot process.

After the reboots, the system will be at the LOADER prompt.

- e. Set the required bootargs:

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

- f. From the LOADER prompt, display the boot menu:

```
boot_ontap menu
```

- g. Again select option “(11) Configure node for external key management” from the boot menu and respond to the prompts as required.

The system boots to the boot menu.

- h. Restore the external key management:

```
security key-manager external restore
```

Creating the network configuration

You must create a network configuration that matches the configuration on the FC nodes. This is because the MetroCluster IP node replays the same configuration when it boots, which means that when node_A_1-IP and node_B_1-IP boot, ONTAP will try to host LIFs on the same ports that were used on node_A_1-FC and node_B_1-FC respectively.

About this task

As you create the network configuration, use the plan made in [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#) to assist you.



Additional configuration may be needed to bring up data LIFs after the MetroCluster IP nodes have been configured.

Steps

1. Verify that all cluster ports are in the appropriate broadcast domain:

The cluster IPspace and cluster broadcast domain are required in order to create cluster LIFs

- a. View the IP spaces:

```
network ipspace show
```

- b. Create IP spaces and assign cluster ports as needed.

[Configuring IPspaces \(cluster administrators only\)](#)

- c. View the broadcast domains:

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

- d. Add any cluster ports to a broadcast domain as needed.

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

- e. Recreate VLANs and interface groups as needed.

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

[Creating a VLAN](#)

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

2. Verify that MTU settings are set correctly for the ports and broadcast domain and make changes using the following commands:

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

```
network port broadcast-domain modify -broadcast-domain bcastdomainname -mtu mtu-value
```

Setting up cluster ports and cluster LIFs

You must set up cluster ports and LIFs. The following steps need to be performed on the site A nodes which were booted up with root aggregates.

Steps

1. Identify the list of LIFs using the desired Cluster port:

```
network interface show -curr-port portname
```

```
network interface show -home-port portname
```

2. For each cluster port, change the home port of any of the LIFs on that port to another port,
 - a. Enter advanced privilege mode and enter “y” when prompted to continue:

```
set priv advanced
```

- b. If the LIF being modified is a data LIF:

```
vserver config override -command "network interface modify -lif lifname -vserver vservername -home-port new-datahomeport"
```

- c. If the LIF is not a data LIF:

```
network interface modify -lif lifname -vserver vservername -home-port new-
```

datahomeport

- d. Revert the modified LIFs to their home port:

```
network interface revert * -vserver vservice_name
```

- e. Verify that there are no LIFs on the cluster port:

```
network interface show -curr-port portname
```

```
network interface show -home-port portname
```

- f. Remove the port from the current broadcast domain:

```
network port broadcast-domain remove-ports -ip-space ipspacename -broadcast-domain bcastdomainname -ports node_name:port_name
```

- g. Add the port to the cluster IPspace and broadcast domain:

```
network port broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster -ports node_name:port_name
```

- h. Verify that the port's role has changed: `network port show`

- i. Repeat these substeps for each cluster port.

- j. Return to admin mode:

```
set priv admin
```

3. Create cluster LIFs on the new cluster ports:

- a. For autoconfiguration using link-local address for cluster LIF, use the following command:

```
network interface create -vserver Cluster -lif cluster_lifname -service -policy default-cluster -home-node a1name -home-port clusterport -auto true
```

- b. To assign static IP address for the cluster LIF, use the following command:

```
network interface create -vserver Cluster -lif cluster_lifname -service -policy default-cluster -home-node a1name -home-port clusterport -address ip-address -netmask netmask -status-admin up
```

Verifying LIF configuration

The node management LIF, cluster management LIF and intercluster LIF will still be present after the storage movement from the old controller. If necessary, you must move LIFs to appropriate ports.

Steps

1. Verify whether the management LIF and cluster management LIFs are on the desired port already:

```
network interface show -service-policy default-management
```

```
network interface show -service-policy default-intercluster
```

If the LIFs are on the desired ports, you can skip the rest of the steps in this task and proceed to the next task.

2. For each node, cluster management, or intercluster LIFs that are not on the desired port, change the home port of any of the LIFs on that port to another port.

- a. Repurpose the desired port by moving any LIFs hosted on desired port to another port:

```
vserver config override -command "network interface modify -lif lifname  
-vserver vservername -home-port new-datahomeport"
```

- b. Revert the modified LIFs to their new home port:

```
vserver config override -command "network interface revert -lif lifname  
-vserver _vservername"
```

- c. If the desired port is not in the right IPspace and broadcast domain, remove the port from the current IPspace and broadcast domain:

```
network port broadcast-domain remove-ports -ipspace current-ip-space  
-broadcast-domain current-broadcast-domain -ports controller-name:current-  
port
```

- d. Move the desired port to the right IPspace and broadcast domain:

```
network port broadcast-domain add-ports -ip-space new-ip-space -broadcast  
-domain new-broadcast-domain -ports controller-name:new-port
```

- e. Verify that the port's role has changed:

```
network port show
```

- f. Repeat these substeps for each port.

3. Move node, cluster management LIFs, and intercluster LIF to the desired port:

- a. Change the LIF's home port:

```
network interface modify -vserver vserver -lif node_mgmt -home-port port  
-home-node homenode
```

- b. Revert the LIF to its new home port:

```
network interface revert -lif node_mgmt -vserver vservername
```

- c. Change the cluster management LIF's home port:

```
network interface modify -vserver vserver -lif cluster-mgmt-LIF-name -home  
-port port -home-node homenode
```

- d. Revert the cluster management LIF to its new home port:

```
network interface revert -lif cluster-mgmt-LIF-name -vserver vservername
```

- e. Change the intercluster LIF's home port:

```
network interface modify -vserver vsverver -lif intercluster-lif-name -home
-node nodename -home-port port
```

f. Revert the intercluster LIF to its new home port:

```
network interface revert -lif intercluster-lif-name -vserver vsververname
```

Bringing up node_A_2-IP and node_B_2-IP

You must bring up and configure the new MetroCluster IP node at each site, creating an HA pair in each site.

Bringing up node_A_2-IP and node_B_2-IP

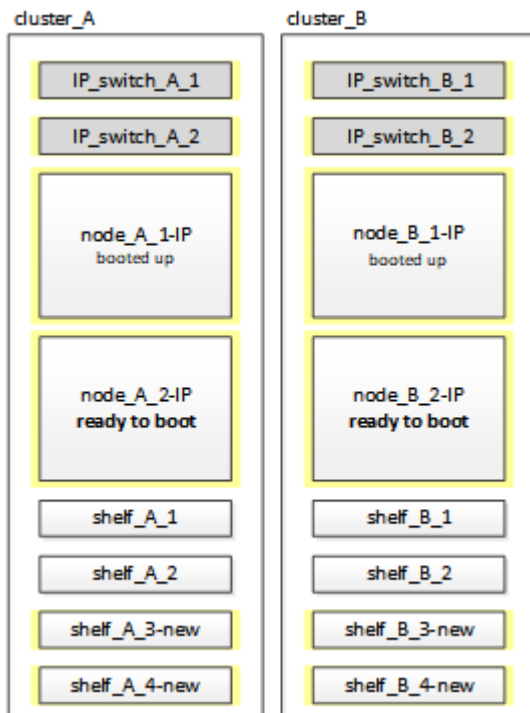
You must boot the new controller modules one at a time using the correct option at the boot menu.

About this task

In these steps, you boot up the two brand new nodes, expanding what had been a two-node configuration into a four-node configuration.

These steps are performed on the following nodes:

- node_A_2-IP
- node_B_2-IP



Steps

1. Boot the new nodes using boot option "9c".

Please choose one of the following:

- (1) Normal Boot.
- (2) Boot without /etc/rc.
- (3) Change password.
- (4) Clean configuration and initialize all disks.
- (5) Maintenance mode boot.
- (6) Update flash from backup config.
- (7) Install new software first.
- (8) Reboot node.
- (9) Configure Advanced Drive Partitioning. Selection (1-9)? 9c

The node initializes and boots to the node setup wizard, similar to the following.

Welcome to node setup

You can enter the following commands at any time:

"help" or "?" - if you want to have a question clarified,

"back" - if you want to change previously answered questions, and

"exit" or "quit" - if you want to quit the setup wizard.

Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a value. .

.
.

If option "9c" does not succeed, take the following steps to avoid possible data loss:

- Do not attempt to run option 9a.
- Physically disconnect the existing shelves that contain data from the original MetroCluster FC configuration (shelf_A_1, shelf_A_2, shelf_B_1, shelf_B_2).
- Contact technical support, referencing the KB article [MetroCluster FC to IP transition - Option 9c Failing](#).

[NetApp Support](#)

2. Enable the AutoSupport tool by following the directions provided by the wizard.
3. Respond to the prompts to configure the node management interface.

Enter the node management interface port: [e0M]:

Enter the node management interface IP address: 10.228.160.229

Enter the node management interface netmask: 225.225.252.0

Enter the node management interface default gateway: 10.228.160.1

4. Verify that the storage failover mode is set to HA:

```
storage failover show -fields mode
```

If the mode is not HA, set it:

```
storage failover modify -mode ha -node localhost
```

You must then reboot the node for the change to take effect.

5. List the ports in the cluster:

```
network port show
```

For complete command syntax, see the man page.

The following example shows the network ports in cluster01:

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

						Speed
(Mbps)						
Node	Port	IPspace	Broadcast Domain	Link	MTU	Admin/Oper
-----	-----	-----	-----	-----	-----	
cluster01-01						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000
cluster01-02						
	e0a	Cluster	Cluster	up	1500	auto/1000
	e0b	Cluster	Cluster	up	1500	auto/1000
	e0c	Default	Default	up	1500	auto/1000
	e0d	Default	Default	up	1500	auto/1000
	e0e	Default	Default	up	1500	auto/1000
	e0f	Default	Default	up	1500	auto/1000

6. Exit the Node Setup wizard:

```
exit
```

7. Log into the admin account using the admin user name.

8. Join the existing cluster using the Cluster Setup wizard.

```

:> cluster setup
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
"help" or "?" - if you want to have a question clarified,
"back" - if you want to change previously answered questions, and "exit"
or "quit" - if you want to quit the cluster setup wizard.
Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
Do you want to create a new cluster or join an existing cluster?
{create, join}:
join

```

9. After you complete the Cluster Setup wizard and it exits, verify that the cluster is active and the node is healthy:

```
cluster show
```

10. Disable disk autoassignment:

```
storage disk option modify -autoassign off -node node_A_2-IP
```

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

If you are using...	Use this command...
Onboard key management	<pre>security key-manager onboard sync</pre> <p>For more information, see Restoring onboard key management encryption keys.</p>
External key management	<pre>security key-manager key query -node node-name</pre> <p>For more information, see Restoring external key management encryption keys.</p>

12. Repeat the above steps on the second new controller module (node_B_2-IP).

Verifying MTU settings

Verify that MTU settings are set correctly for the ports and broadcast domain and make changes.

Steps

1. Check the MTU size used in the cluster broadcast domain:

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

2. If necessary, update the MTU size as needed:


```
network port broadcast-domain modify -broadcast-domain bcast-domain-name -mtu
mtu-size
```

Configuring intercluster LIFs

Configure the intercluster LIFs required for cluster peering.

This task must be performed on both of the new nodes, `node_A_2-IP` and `node_B_2-IP`.

Step

1. Configure the intercluster LIFs. See [Configuring intercluster LIFs](#)

Verifying cluster peering

Verify that `cluster_A` and `cluster_B` are peered and nodes on each cluster can communicate with each other.

Steps

1. Verify the cluster peering relationship:

```
cluster peer health show
```

```
cluster01::> cluster peer health show
Node          cluster-Name          Node-Name
          Ping-Status          RDB-Health Cluster-Health Avail...
-----
node_A_1-IP
          cluster_B          node_B_1-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          node_B_2-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
node_A_2-IP
          cluster_B          node_B_1-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
          node_B_2-IP
          Data: interface_reachable
          ICMP: interface_reachable true          true          true
```

2. Ping to check that the peer addresses are reachable:

```
cluster peer ping -originating-node local-node -destination-cluster remote-
cluster-name
```

Configuring the new nodes and completing transition

With the new nodes added, you must complete the transition steps and configure the MetroCluster IP nodes.

Configuring the MetroCluster IP nodes and disabling transition

You must implement the MetroCluster IP connections, refresh the MetroCluster configuration, and disable transition mode.

- 1. Form the new nodes into a DR group by issuing the following commands from controller node_A_1-IP:

```
metrocluster configuration-settings dr-group create -partner-cluster peer-cluster-name -local-node local-controller-name -remote-node remote-controller-name
```

```
metrocluster configuration-settings dr-group show
```

- 2. Create MetroCluster IP interfaces (node_A_1-IP, node_A_2-IP, node_B_1-IP, node_B_2-IP) — two interfaces need to be created per controller; eight interfaces in total:

```
metrocluster configuration-settings interface create -cluster-name cluster-name -home-node controller-name -home-port port -address ip-address -netmask netmask -vlan-id vlan-id`metrocluster configuration-settings interface show
```



Beginning with ONTAP 9.9.1, if you are using a layer 3 configuration, you must also specify the `-gateway` parameter when creating MetroCluster IP interfaces. Refer to [Considerations for layer 3 wide-area networks](#).

The `-vlan-id` parameter is required only if you are not using the default VLAN IDs. Only certain systems support non-default VLAN IDs.



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

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

AFF platforms	FAS platforms
---------------	---------------

<ul style="list-style-type: none"> • AFF A220 • AFF A250 • AFF A400 	<ul style="list-style-type: none"> • FAS2750 • FAS500f • FAS8300 • FAS8700
--	--

3. Perform the MetroCluster connect operation from controller node_A_1-IP to connect the MetroCluster sites — this operation can take a few minutes to complete:

```
metrocluster configuration-settings connection connect
```

4. Verify that the remote cluster disks are visible from each controller via the iSCSI connections:

```
disk show
```

You should see the remote disks belonging to the other nodes in the configuration.

5. Mirror the root aggregate for node_A_1-IP and node_B_1-IP:

```
aggregate mirror -aggregate root-aggr
```

6. Assign disks for node_A_2-IP and node_B_2-IP.

Pool 1 disk assignments were already made for node_A_1-IP and node_B_1-IP when the `boot_after_mcc_transition` command was issued at the boot menu.

- a. Issue the following commands on node_A_2-IP:

```
disk assign disk1disk2disk3 ... diskn -sysid node_B_2-IP-controller-sysid
-pool 1 -force
```

- b. Issue the following commands on node_B_2-IP:

```
disk assign disk1disk2disk3 ... diskn -sysid node_A_2-IP-controller-sysid
-pool 1 -force
```

7. Confirm ownership has been updated for the remote disks:

```
disk show
```

8. If necessary, refresh the ownership information using the following commands:

- a. Go to advanced privilege mode and enter `y` when prompted to continue:

```
set priv advanced
```

- b. Refresh disk ownership:

```
disk refresh-ownership controller-name
```

- c. Return to admin mode:

```
set priv admin
```

9. Mirror the root aggregates for node_A_2-IP and node_B_2-IP:

```
aggregate mirror -aggregate root-aggr
```

10. Verify that the aggregate re-synchronization has completed for root and data aggregates:

```
aggr show`aggr plex show
```

The resync can take some time but must complete before proceeding with the following steps.

11. Refresh the MetroCluster configuration to incorporate the new nodes:

a. Go to advanced privilege mode and enter y when prompted to continue:

```
set priv advanced
```

b. Refresh the configuration:

If you have configured...	Issue this command...
A single aggregate in each cluster:	metrocluster configure -refresh true -allow-with-one-aggregate true
More than a single aggregate in each cluster	metrocluster configure -refresh true

c. Return to admin mode:

```
set priv admin
```

12. Disable MetroCluster transition mode:

a. Enter advanced privilege mode and enter “y” when prompted to continue:

```
set priv advanced
```

b. Disable transition mode:

```
metrocluster transition disable
```

c. Return to admin mode:

```
set priv admin
```

Setting up data LIFs on the new nodes

You must configure data LIFs on the new nodes, node_A_2-IP and node_B_2-IP.

You must add any new ports available on new controllers to a broadcast domain if not already assigned to one. If required, create VLANs or interface groups on the new ports. See [Network management](#)

1. Identify the current port usage and broadcast domains:

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

2. Add ports to broadcast domains and VLANs as necessary.

- a. View the IP spaces:

```
network ipspace show
```

- b. Create IP spaces and assign data ports as needed.

[Configuring IPspaces \(cluster administrators only\)](#)

- c. View the broadcast domains:

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

- d. Add any data ports to a broadcast domain as needed.

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

- e. Recreate VLANs and interface groups as needed.

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

[Creating a VLAN](#)

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

3. Verify that the LIFs are hosted on the appropriate node and ports on the MetroCluster IP nodes (including the SVM with -mc vserver) as needed.

See the information gathered in [Creating the network configuration](#).

- a. Check the home port of the LIFs:

```
network interface show -field home-port
```

- b. If necessary, modify the LIF configuration:

```
vserver config override -command "network interface modify -vserver  
vserver_name -home-port active_port_after_upgrade -lif lif_name -home- node  
new_node_name"
```

- c. Revert the LIFs to their home ports:

```
network interface revert * -vserver vserver_name
```

Bringing up the SVMs

Due to the changes if LIF configuration, you must restart the SVMs on the new nodes.

Steps

1. Check the state of the SVMs:

```
metrocluster vserver show
```

2. Restart the SVMs on cluster_A that do not have an “-mc” suffix:

```
vserver start -vserver svm-name -force true
```

3. Repeat the previous steps on the partner cluster.

4. Check that all SVMs are in a healthy state:

```
metrocluster vserver show
```

5. Verify that all data LIFs are online:

```
network interface show
```

Moving a system volume to the new nodes

To improve resiliency, a system volume should be moved from controller node_A_1-IP to controller node_A_2-IP, and also from node_B_1-IP to node_B_2-IP. You must create a mirrored aggregate on the destination node for the system volume.

About this task

System volumes have the name form “MDV_CRS_*_A” or “MDV_CRS_*_B.” The designations “_A” and “_B” are unrelated to the site_A and site_B references used throughout this section; e.g., MDV_CRS_*_A is not associated with site_A.

Steps

1. Assign at least three pool 0 and three pool 1 disks each for controllers node_A_2-IP and node_B_2-IP as needed.
2. Enable disk auto-assignment.
3. Move the _B system volume from node_A_1-IP to node_A_2-IP using the following steps from site_A.
 - a. Create a mirrored aggregate on controller node_A_2-IP to hold the system volume:

```
aggr create -aggregate new_node_A_2-IP_aggr -diskcount 10 -mirror true -node  
nodename_node_A_2-IP
```

```
aggr show
```

The mirrored aggregate requires five pool 0 and five pool 1 spare disks owned by controller node_A_2-IP.

The advanced option, “-force-small-aggregate true” can be used to limit disk use to 3 pool 0 and 3 pool 1 disks, if disks are in short supply.

- b. List the system volumes associated with the admin SVM:

```
vserver show
```

```
volume show -vserver admin-vserver-name
```

You should identify volumes contained by aggregates owned by site_A. The site_B system volumes will also be shown.

4. Move the MDV_CRS_*_B system volume for site_A to the mirrored aggregate created on controller node_A_2-IP

- a. Check for possible destination aggregates:

```
volume move target-aggr show -vserver admin-vserver-name -volume
system_vol_MDV_B
```

The newly created aggregate on node_A_2-IP should be listed.

- b. Move the volume to the newly created aggregate on node_A_2-IP:

```
set advanced
```

```
volume move start -vserver admin-vserver -volume system_vol_MDV_B
-destination-aggregate new_node_A_2-IP_aggr -cutover-window 40
```

- c. Check status for the move operation:

```
volume move show -vserver admin-vserver-name -volume system_vol_MDV_B
```

- d. When the move operation complete, verify that the MDV_CRS_*_B system is contained by the new aggregate on node_A_2-IP:

```
set admin
```

```
volume show -vserver admin-vserver
```

5. Repeat the above steps on site_B (node_B_1-IP and node_B_2-IP).

Returning the system to normal operation

You must perform final configuration steps and return the MetroCluster configuration to normal operation.

Verifying MetroCluster operation and assigning drives after transition

You must verify that the MetroCluster is operating correctly and assign drives to the second pair of new nodes (node_A_2-IP and node_B_2-IP).

1. Confirm that the MetroCluster configuration-type is IP-fabric: `metrocluster show`
2. Perform a MetroCluster check.
 - a. Issue the following command: `metrocluster check run`
 - b. Display the results of the MetroCluster check: `metrocluster check show`
3. Confirm that the DR group with the MetroCluster IP nodes is configured: `metrocluster node show`
4. Create and mirror additional data aggregates for controllers node_A_2-IP and node_B_2-IP at each site as needed.

Installing licenses for the new controller module

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

For detailed information about licensing, see the knowledgebase article 3013749: Data ONTAP 8.2 Licensing Overview and References on the NetApp Support Site and the *System Administration Reference*.

1. If necessary, obtain license keys for the new node on the NetApp Support Site in the My Support section under Software licenses.

For further information on license replacements, see the Knowledge Base article [Post Motherboard Replacement Process to update Licensing on a AFF/FAS system](#).

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

The license_key is 28 digits in length.

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

Completing configuration of the nodes

There are miscellaneous configuration steps that can be performed prior to completing the procedures. Some of these steps are optional.

1. Configure the service processor: `system service-processor network modify`
2. Set up autosupport on the new nodes: `system node autosupport modify`
3. The controllers can be optionally renamed as part of the transition. The following command is used to rename a controller: `system node rename -node <old-name> -newname <new-name>`

The renaming operation can take a few minutes to complete. Confirm that any name changes have propagated to each node prior to continuing with other steps using the `system show -fields node` command.

4. Configure a monitoring service as desired.

[Considerations for Mediator](#)

`xref:./transition/./install-ip/concept_mediator_requirements.html`

[Tiebreaker Software installation and configuration](#)

Sending a custom AutoSupport message after maintenance

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

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

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

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.