



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

ONTAP MetroCluster

NetApp
September 06, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-metrocluster/transition/concept_nondisruptively_transitioning_from_a_four_node_mcc_fc_to_a_mcc_ip_configuration.html on September 06, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Transition nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later) . . 1
 - Transitioning nondisruptively from a MetroCluster FC to a MetroCluster IP configuration (ONTAP 9.8 and later) 1
 - Prepare for transition from a MetroCluster FC to a MetroCluster IP configuration 1
 - Transition from MetroCluster FC to MetroCluster IP configurations. 12
 - Sending a custom AutoSupport message after maintenance 62
 - Restoring Tiebreaker or Mediator monitoring. 62

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

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

You can perform nondisruptive transitions of workloads and data from an existing MetroCluster FC configuration to a new MetroCluster IP configuration.

Beginning with ONTAP 9.13.1, this procedure is supported in MetroCluster IP configurations in which the MetroCluster and the drive shelves are connected to the same IP switches (a shared storage switch configuration).

Beginning with ONTAP 9.13.1, you can perform a nondisruptive transition of workloads and data from an existing eight-node MetroCluster FC configuration to a new MetroCluster IP configuration.

Beginning with ONTAP 9.8, you can perform a nondisruptive transition of workloads and data from an existing four-node MetroCluster FC configuration to a new MetroCluster IP configuration.

- This procedure is nondisruptive.

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

- This procedure applies only to four-node and eight-node MetroCluster FC configurations.

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

- This procedure describes the steps required to transition one four-node FC DR group. If you have an eight-node configuration (two FC DR groups), you must repeat the entire procedure for each FC DR group.
- You must meet all requirements and follow all steps in the procedure.

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

Enable console logging

Enable console logging on your devices before performing this task.

NetApp strongly recommends that you enable console logging on the devices that you are using and take the following actions when performing this procedure:

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

See the Knowledge Base article [How to suppress automatic case creation during scheduled maintenance windows](#).

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

Requirements for nondisruptive FC-to-IP transition

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

- If you have an eight-node configuration, all nodes must be running ONTAP 9.13.1 or later.
- If you have a four-node configuration, 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](#)

- It must support a switched cluster configuration.

[NetApp Hardware Universe](#)

- 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](#)

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 and FC-to-SAS bridges, 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.

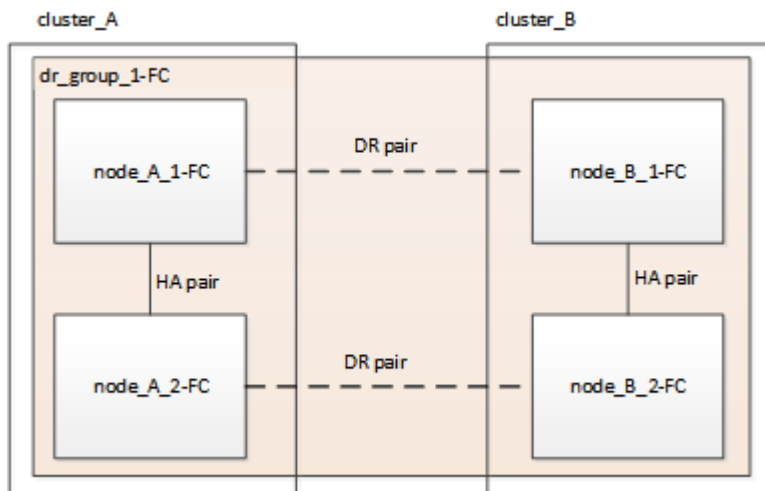
Workflow for nondisruptive MetroCluster transition

You must follow the specific workflow to ensure a successful nondisruptive transition. Choose the workflow for your configuration:


- [Four-node FC configuration transition workflow](#)
- [Eight-node FC configuration transition workflow](#)

Four-node FC configuration transition workflow


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.




Eight-node FC configuration transition workflow

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




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




Data is transferred from DR_group_1-FC to DR_group_1-IP, and then the old nodes and their storage are removed from the configuration and decommissioned.




If you want to transition from an eight-node FC configuration to a four-node IP configuration, you must transition all the data in DR_group_1-FC and DR_group_2-FC to the new IP DR group (DR_group_1-IP). You can then decommission both FC DR groups. After the FC DR groups have been removed, the process ends with a four-node MetroCluster IP configuration.



Add the remaining MetroCluster IP nodes to the existing MetroCluster configuration. Repeat the process to transfer data from the DR_group_2-FC nodes to the DR_group_2-IP nodes.




After removing DR_group_2-FC, the process ends with an eight-node MetroCluster IP configuration.



Transition process workflow

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 NetApp-provided switches.

- The use of MetroCluster-compliant switches (switches that are not validated and provided by NetApp) 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.

See the *NetApp Hardware Universe* for information on platform model and switch support.

[NetApp Hardware Universe](#)

Switchover, healing, and switchback operations during nondisruptive 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 those used by a MetroCluster FC configuration.
- In the MetroCluster IP workflow, the switchover, healing, and switchback steps are those 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.

For information on the MetroCluster FC and IP workflows for switchover, healing, and switchback, see [Understanding MetroCluster data protection and disaster recovery](#).



Automatic unplanned switchover is not available during the transition process.

Stage of transition	Negotiated switchover uses this workflow...	Unplanned switchover uses this workflow...

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. • The <code>heal aggregate</code> and <code>heal root</code> 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.

- OnCommand 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
Switches	<p>Initial switches (if fabric-attached configuration):</p> <ul style="list-style-type: none"> • switch_A_1-FC • switch_A_2-FC <p>MetroCluster IP switches:</p> <ul style="list-style-type: none"> • switch_A_1-IP • switch_A_2-IP 	<p>Initial switches (if fabric-attached configuration):</p> <ul style="list-style-type: none"> • switch_B_1-FC • switch_B_2-FC <p>MetroCluster IP switches:</p> <ul style="list-style-type: none"> • switch_B_1-IP • switch_B_2-IP

Transition from MetroCluster FC to MetroCluster IP configurations

Verifying the health of the MetroCluster configuration

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

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 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::>
```

3. 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::>
```

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 "up/up" for "Status Admin/Oper".

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

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----
-----	-----				
Cluster					
	node_A-1_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::>
```

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.

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.

This task is required if you are using new switches.

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

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](#)

- [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 *MetroCluster IP Installation and Configuration*.

[MetroCluster IP installation and configuration](#)

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

Move the local cluster connections

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

Move 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.

You must perform this task on both MetroCluster sites.

Which connections to move

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, 3232C, or 9336C-FX2	Any	5	Local cluster interface on FC node
		6	Local cluster interface on FC node
Broadcom BES-53248	FAS500f/A250	1 - 6	Local cluster interface on FC node
	FAS8200/A300	3, 4, 9, 10, 11, 12	Local cluster interface on FC node
	FAS8300/A400/FAS8700	1 - 6	Local cluster interface on FC node

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.

1. Move the MetroCluster FC node group A cluster connections to the new IP switches.

Use the ports described in [Which connections to move](#).

- 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) Admin/Oper	Health 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) 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::*>
```

3. Verify that your inter-site Inter-Switch Links (ISLs) are up and the port-channels are operational:

```
show interface brief
```

In the following example, ISL ports “Eth1/15” to “Eth1/20” are configured as “Po10” for the remote site link and “Eth1/7” to “Eth1/8” are configured as “Po1” for the local cluster ISL. The state of “Eth1/15” to “Eth1/20”, “Eth1/7” to “Eth1/8”, “Po10”, and “Po1” should be 'up'.

```
IP_switch_A_1# show interface brief
```

Port	VRF	Status	IP Address	Speed	MTU
mgmt0	--	up	100.10.200.20	1000	1500

Ethernet Port Interface	VLAN	Type	Mode	Status	Reason	Speed
					Ch #	

...

```

Eth1/7      1      eth  trunk  up      none      100G(D)
1
Eth1/8      1      eth  trunk  up      none      100G(D)
1
...

Eth1/15     1      eth  trunk  up      none      100G(D)
10
Eth1/16     1      eth  trunk  up      none      100G(D)
10
Eth1/17     1      eth  trunk  up      none      100G(D)
10
Eth1/18     1      eth  trunk  up      none      100G(D)
10
Eth1/19     1      eth  trunk  up      none      100G(D)
10
Eth1/20     1      eth  trunk  up      none      100G(D)
10

-----
-----
Port-channel VLAN  Type Mode  Status  Reason      Speed      Protocol
Interface
-----
-----
Po1          1      eth  trunk  up      none      a-100G(D)  lacp
Po10         1      eth  trunk  up      none      a-100G(D)  lacp
Po11         1      eth  trunk  down    No operational  auto(D)    lacp
members
IP_switch_A_1#

```

4. 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::*>
```

5. 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.
6. 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.

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

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.

[AFF A320 systems: Installation and setup](#)

[AFF A220/FAS2700 Systems Installation and Setup Instructions](#)

[AFF A800 Systems Installation and Setup Instructions](#)

[AFF A300 Systems Installation and Setup Instructions](#)

[FAS8200 Systems Installation and Setup Instructions](#)

2. Reconfigure switch_A_1_IP using RCF files generated for your platform combination and transition.

Follow the steps in the procedure for your switch vendor from *MetroCluster IP Installation and Configuration*:

[MetroCluster IP installation and configuration](#)

- 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 [Which connections to move](#).

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) Admin/Oper	Health 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) 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::*>
```

5. Verify that all interfaces are on their home port:

```
network interface show -vserver Cluster
```



```
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::*>
```

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.

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) Admin/Oper	Health 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) 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::*>
```

2. Verify that all interfaces are on their home port:

```
network interface show -vserver Cluster
```

This might 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
```

	Logical	Status	Network	Current	
Current Is					
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::*>
```

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				
	e0b	Switch-A-4-IP	0/5/1	N3K-
C3232C				
Node-A-2-FC				
	/cdp			
	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 ID	Local Intrfce	Hldtme	Capability	Platform	Port
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 (FDO220329A4)	Eth1/7	175	R S I s	N3K-C3232C	Eth1/7
Switch-A-4-IP (FDO220329A4)	Eth1/8	175	R S I s	N3K-C3232C	Eth1/8
Switch-B-3-IP (FDO220329B3)	Eth1/20	173	R S I s	N3K-C3232C	
Eth1/20					
Switch-B-3-IP (FDO220329B3)	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 that 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 ID	Local Intrfce	Hldtme	Capability	Platform	Port
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.

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.

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 [Cabling the IP switches](#).

3. Configure the MetroCluster IP nodes using the following sections:
 - 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 \(ONTAP 9.4 and later\)](#)
4. From Maintenance mode, issue the halt command to exit Maintenance mode, and then issue the boot_ontap command to boot the system and get to cluster setup.

Do not complete the cluster wizard or node wizard at this time.

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 NetApp technical support that maintenance is underway. Informing technical support that maintenance is underway prevents them from opening a case on the assumption that a disruption has occurred.

About this task

This task must be performed on each MetroCluster site.

Steps

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

```
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
```

2. 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.

1. Enable transition:
 - a. Change to the advanced privilege level:

```
set -privilege advanced
```

b. Enable transition mode:

```
metrocluster transition enable -transition-mode non-disruptive
```



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.



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.



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.

Steps

1. Add the MetroCluster IP nodes to the existing MetroCluster configuration.
 - a. Join the first MetroCluster IP node (node_A_3-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 NetApp Technical Support. To disable this feature, enter `autosupport modify -support disable` within 24 hours.

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

Type yes to confirm and continue {yes}: yes

Enter the node management interface port [e0M]:
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_4-IP) to the existing MetroCluster FC configuration.
2. Repeat these steps to join node_B_3-IP and node_B_4-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.

Steps

1. On the new MetroCluster IP nodes, [configure the intercluster LIFs](#).
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
2	cluster_A	node_A_3-IP	node_B_3-IP
		node_A_4-IP	node_B_4-IP
	cluster_B	node_B_3-IP	node_A_3-IP
		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_3-IP	ready to configure	-	-
		node_A_4-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_3-IP	ready to configure	-	-
		node_B_4-IP	ready to configure	-	-

5. Configure the MetroCluster IP interfaces for the newly joined MetroCluster IP nodes:

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

See [Configuring and connecting the MetroCluster IP interfaces](#) for considerations when configuring the IP interfaces.



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
Config
Group Cluster Node      Network Address Netmask      Gateway
State
-----
-----
2      cluster_A
      node_A_3-IP
      Home Port: e1a
      172.17.26.10      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.10      255.255.255.0      -
completed
      node_A_4-IP
      Home Port: e1a
      172.17.26.11      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.11      255.255.255.0      -
completed
      cluster_B
      node_B_3-IP
      Home Port: e1a
      172.17.26.13      255.255.255.0      -
completed
      Home Port: e1b
      172.17.27.13      255.255.255.0      -
completed
      node_B_3-IP
      Home Port: e1a
      172.17.26.12      255.255.255.0      -
completed
      Home Port: e1b
      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
```




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
2	cluster_A	node_A_3-IP**	Home Port: e1a 172.17.26.10	172.17.26.11	HA Partner
completed			Home Port: e1a 172.17.26.10	172.17.26.12	DR Partner
completed			Home Port: e1a 172.17.26.10	172.17.26.13	DR Auxiliary
completed			Home Port: e1b 172.17.27.10	172.17.27.11	HA Partner
completed			Home Port: e1b 172.17.27.10	172.17.27.12	DR Partner
completed			Home Port: e1b 172.17.27.10	172.17.27.13	DR Auxiliary
		node_A_4-IP	Home Port: e1a 172.17.26.11	172.17.26.10	HA Partner
completed			Home Port: e1a 172.17.26.11	172.17.26.13	DR Partner
completed			Home Port: e1a		

```

completed          172.17.26.11      172.17.26.12      DR Auxiliary
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                  Source          Destination
Group Cluster Node  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 Owner	Usable Size	Shelf	Bay	Disk Type	Container Type	Container Name
-----	-----	-----	-----	-----	-----	-----
1.10.4	-	10	4	SAS	remote	-
node_B_2						
1.10.13	-	10	13	SAS	remote	-
node_B_2						
1.10.14	-	10	14	SAS	remote	-
node_B_1						
1.10.15	-	10	15	SAS	remote	-
node_B_1						
1.10.16	-	10	16	SAS	remote	-
node_B_1						
1.10.18	-	10	18	SAS	remote	-
node_B_2						
...						
2.20.0	546.9GB	20	0	SAS	aggregate	aggr0_rha1_a1
node_a_1						
2.20.3	546.9GB	20	3	SAS	aggregate	aggr0_rha1_a2
node_a_2						
2.20.5	546.9GB	20	5	SAS	aggregate	rha1_a1_aggr1
node_a_1						
2.20.6	546.9GB	20	6	SAS	aggregate	rha1_a1_aggr1
node_a_1						
2.20.7	546.9GB	20	7	SAS	aggregate	rha1_a2_aggr1
node_a_2						
2.20.10	546.9GB	20	10	SAS	aggregate	rha1_a1_aggr1
node_a_1						
...						
43 entries were displayed.						

```
cluster_A::>
```



On systems configured for Advanced Drive Partitioning (ADP), the container type is "shared" rather than "remote" as shown in the example output.

10. Mirror the root aggregates:

```
storage aggregate mirror -aggregate aggr0_node_A_3_IP
```



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)

Physical Size	Position	Disk	Type	Usable Size
-----	-----	-----	-----	-----
-----	dparity	4.20.0	SAS	-
-	parity	4.20.3	SAS	-
-	data	4.20.1	SAS	546.9GB
558.9GB				

Aggregate capacity available for volume use would be 467.6GB.

Do you want to continue? {y|n}: y

```
cluster_A::>
```

11. Verify that the root aggregates are mirrored:

```
storage aggregate show
```

```
cluster_A::> aggr show
```

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

Steps

1. Configure the MetroCluster depending on whether it has a single or multiple data aggregates:

If your MetroCluster configuration has...	Then do this...
---	-----------------

Multiple data aggregates	<p>From any node's prompt, configure MetroCluster:</p> <pre>metrocluster configure <node-name></pre> <div>  <p>You must run <code>metrocluster configure</code> and not <code>metrocluster configure -refresh true</code></p> </div>
A single mirrored data aggregate	<p>a. From any node's prompt, change to the advanced privilege level:</p> <pre>set -privilege advanced</pre> <p>You must respond with <code>y</code> when you are prompted to continue into advanced mode and you see the advanced mode prompt (<code>*></code>).</p> <p>b. Configure the MetroCluster with the <code>-allow-with-one-aggregate true</code> parameter:</p> <pre>metrocluster configure -allow-with-one-aggregate true -node-name <node-name></pre> <p>c. Return to the admin privilege level:</p> <pre>set -privilege admin</pre>



The best practice is to have multiple mirrored data aggregates. When there is only one mirrored aggregate, there is less protection because the metadata volumes are located on the same aggregate rather than on separate aggregates.

2. Verify that the nodes are added to their DR group:

```
metrocluster node show
```

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

DR	Configuration	DR
Group Cluster Node	State	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::>
```

3. Create mirrored data aggregates on each of the new MetroCluster nodes:

```
storage aggregate create -aggregate aggregate-name -node node-name -diskcount
no-of-disks -mirror true
```



You must create at least one mirrored data aggregate per site. It is recommended to have two mirrored data aggregates per site on MetroCluster IP nodes to host the MDV volumes, however a single aggregate per site is supported (but not recommended). It is acceptable that one site of the MetroCluster has a single mirrored data aggregate and the other site has more than one mirrored data aggregate.

The following example shows the creation of an aggregate on node_A_3-IP.

```
cluster_A::> storage aggregate create -aggregate data_a3 -node node_A_3-
IP -diskcount 10 -mirror t
```

Info: The layout for aggregate "data_a3" on node "node_A_3-IP" would be:

First Plex

RAID Group rg0, 5 disks (block checksum, raid_dp)

Usable

Physical

Position

Disk

Type

Size


```

Size
-----
-----
-      dparity    5.10.15          SAS          -
-      parity     5.10.16          SAS          -
-      data       5.10.17          SAS          546.9GB
547.1GB
-      data       5.10.18          SAS          546.9GB
558.9GB
-      data       5.10.19          SAS          546.9GB
558.9GB

```

Second Plex

RAID Group rg0, 5 disks (block checksum, raid_dp)

```

Usable
Physical
Position  Disk          Type          Size
Size
-----
-----
-      dparity    4.20.17          SAS          -
-      parity     4.20.14          SAS          -
-      data       4.20.18          SAS          546.9GB
547.1GB
-      data       4.20.19          SAS          546.9GB
547.1GB
-      data       4.20.16          SAS          546.9GB
547.1GB

```

Aggregate capacity available for volume use would be 1.37TB.

Do you want to continue? {y|n}: y

[Job 440] Job succeeded: DONE

cluster_A::>

4. Verify that all nodes in the cluster are healthy:

```
cluster show
```

The output should display `true` for the `health` field for all nodes.

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

```
storage failover show
```

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

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

6. Confirm that all disks attached to the newly-joined MetroCluster IP nodes are present:

```
disk show
```

7. Verify the health of the MetroCluster configuration by running the following commands:

- metrocluster check run
- metrocluster check show
- metrocluster interconnect mirror show
- metrocluster interconnect adapter show

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

- Display the volumes to identify the MDV volumes:



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

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

```

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

cluster_A  MDV_CRS_2c78e009ff5611e9b0f300a0985ef8c4_A
            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.

Before you begin

You should create new SAN LIFs on the destination or IP nodes and connect hosts prior to moving volumes to new the new aggregates.

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

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

2. Move the data volumes to aggregates on the new controllers, one volume at a time.

Use the procedure in [Creating an aggregate and moving volumes to the new nodes](#).

3. Create SAN LIFs on the recently added nodes.

Use the following procedure in [Updating LUN paths for the new nodes](#).

4. 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 procedure in [Adding node-locked licenses](#).

5. Migrate the data LIFs.

Use the procedure in [Moving non-SAN data LIFs and cluster management LIFs to the new nodes](#) but do **not** perform the last two steps to migrate cluster management LIFs.



- You cannot migrate a LIF that is used for copy-offload operations with VMware vStorage APIs for Array Integration (VAAI).
- After you complete the transition of your MetroCluster nodes from FC to IP, you might need to move your iSCSI host connections to the new nodes, see [Moving Linux iSCSI hosts from MetroCluster FC to MetroCluster IP nodes](#).

Removing the MetroCluster FC controllers

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

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

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

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

- b. Repeat the command on the partner cluster.

2. 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`.



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 Status	Size	Available	Used%	State	#Vols	Nodes	RAID
-----	-----	-----	-----	-----	-----	-----	-----
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 data aggregates offline, and then delete them:

- a. Take the aggregate offline: `storage aggregate offline -aggregate aggregate-name`

The following example shows the aggregate `aggr_data_a1` being taken offline:

```

cluster_B::> storage aggregate offline -aggregate aggr_data_a1

Aggregate offline successful on aggregate: aggr_data_a1

```

- b. Delete the aggregate: `storage aggregate delete -aggregate aggregate-name`

You can destroy the plex when prompted.

The following example shows the aggregate `aggr_data_a1` being deleted.

```
cluster_B::> storage aggregate delete -aggregate aggr_data_a1
Warning: Are you sure you want to destroy aggregate "aggr_data_a1"?
{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 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_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:

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

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

```
cluster_A::> cluster peer show
Peer Cluster Name      Cluster Serial Number Availability
Authentication
-----
cluster_B              1-80-000011              Unavailable      absent
```

- i. Modify the cluster_A peer IP address:

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

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

```
cluster_B::> cluster peer show
Peer Cluster Name      Cluster Serial Number Availability
Authentication
-----
cluster_A              1-80-000011          Unavailable      absent
```

i. Modify the cluster_B peer IP address:

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

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

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

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

Example for cluster_A:

```
cluster_A::> cluster peer show -instance

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

cluster_A::>
```

Example for cluster_B

```
cluster_B::> cluster peer show -instance

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

cluster_B::>
```

10. 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::>
```

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

You must perform this step on both clusters.



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::>
```

12. 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::>
```

13. 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::>

```

14. Power down the MetroCluster FC controller modules and storage shelves.
15. 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.

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
                                Takeover
Node        Partner             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`
- c. Return to the admin privilege level: `set -privilege admin`

```
cluster_A::*> metrocluster transition disable

cluster_A::*>
```

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::>
```

6. If you have an eight-node configuration, you must repeat the entire procedure starting from [Prepare for transition from a MetroCluster FC to a MetroCluster IP configuration](#) for each of the FC DR groups.

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.

Restoring Tiebreaker or Mediator monitoring

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

- 1. Use the appropriate procedure for your configuration.

If you are using...	Use this procedure
Tiebreaker	Adding MetroCluster configurations
Mediator	Configuring the ONTAP Mediator service from a MetroCluster IP configuration

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.