



Upgrading controller hardware in a pair of nodes running clustered Data ONTAP 8.3 by moving storage

To upgrade the controller hardware in a pair of nodes running clustered Data ONTAP 8.3, you must read and understand the guidelines, prepare both pairs of nodes, and then replace the old node pair with the new node pair. During the upgrade, you can transfer storage from the original system in four ways: transfer external disk shelves, transfer internal disk drives from a FAS2220A or FAS32xx system, transfer a connected storage third-party storage array, or convert a FAS2240A or FAS25xx system to a disk shelf and attach it to the new system.

About this task

- This is a disruptive procedure.
If you want to perform a nondisruptive upgrade of controller hardware in a pair of nodes, see the procedure *Using Aggregate Relocation to upgrade controller hardware on a pair of nodes running clustered Data ONTAP 8.3* or the procedure *Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP by moving volumes*.
- This procedure is written with the following assumptions:
 - You are upgrading a pair of nodes running clustered Data ONTAP 8.3 to a new pair of nodes running clustered Data ONTAP 8.3.
If you want to upgrade a pair of nodes not currently running Data ONTAP 8.3, follow the instructions in the *Ensuring hardware and software compatibility between original and new nodes* section to upgrade Data ONTAP.
 - You are reusing the IP addresses, netmasks, and gateways of the original nodes on the new nodes.
- After the procedure, you will use the root volume of the original system.
- If you are replacing an individual component, see the field-replaceable unit (FRU) flyer for that component on the NetApp Support Site at mysupport.netapp.com.
- This procedure uses the term *boot environment prompt* to refer to the prompt on a node from which you can perform certain tasks, such as rebooting the node and printing or setting environmental variables.
The prompt is shown in the following example:

```
LOADER>
```

This prompt is often referred to as the *boot loader prompt*.

Note: If you need help with any part of this procedure, contact technical support.

This procedure applies to FAS systems, V-Series systems, and systems with FlexArray Virtualization Software.

Note: Most Data ONTAP platforms released before Data ONTAP 8.2.1 were released as separate FAS and V-Series hardware systems (for example, a FAS6280 and a V6280). Only the V-Series systems (a V or GF prefix) could attach to storage arrays. Starting in Data ONTAP 8.2.1, only one hardware system is being released for new platforms. These new platforms, which have only a FAS prefix, can attach to storage arrays if the required license is installed. These new platforms are the FAS80xx and FAS25xx systems.

This document uses the term “systems with FlexArray Virtualization Software” to refer to systems belonging to these new platforms and the term “V-Series system” to refer to the separate hardware systems that can attach to storage arrays.

Steps

1. [Guidelines for upgrading a pair of nodes](#) on page 2
2. [Required tools and documentation](#) on page 3
3. [Preparing to upgrade the node pair](#) on page 4
4. [Upgrading the node pair](#) on page 12
5. [Performing post-upgrade tasks](#) on page 27

Guidelines for upgrading a pair of nodes

You need to be aware of several issues before you upgrade the node pair.

Supported upgrade paths

You can upgrade a pair of nodes as follows:

Original nodes	New nodes
FAS22xxA See the note below the table.	FAS25xx, 3220, 3250, 6220, 6250, 6290, FAS80xx
FAS2520	FAS25xx, 3220, 3250, 6220, 6250, 6290, FAS80xx
FAS255x	FAS25xx, 3220, 3250, 6220, 6250, 6290, FAS80xx
3220	3250, 6220, 6250, 6290, FAS80xx
3240, 3270	3250, 6220, 6250, 6290, FAS80xx
3250	6220, 6250, 6290, FAS80xx
6210	6220, FAS8060, FAS8080 EX
6220	FAS8060, FAS8080 EX
6240	6250, 6290, FAS8060 with I/O module, FAS8080 EX
6280	FAS8080 EX

Note: If your FAS80xx controllers are running Data ONTAP 8.3 or later and one or both are All-Flash FAS models, make sure that both controllers have the same All-Flash Optimized personality set:

```
system node show -instance node_name
```

Both nodes must either have the personality enabled or disabled; you cannot combine a node with the All-Flash Optimized personality enabled with a node that does not have the personality enabled in the same HA pair. If the personalities are different, refer to [KB Article 1015157](#) in the NetApp Knowledge Base for instructions on how to sync node personality.

Note: If the root aggregate on FAS22xx systems is on partitioned disks in clustered Data ONTAP 8.3 you will not be able to use this procedure. You can upgrade by using the procedure *Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP by moving volumes*. After the data is removed, you can delete the existing partitions and remove ownership of the disks and then the FAS22xx can be repurposed as a data shelf.

Note: The system being upgraded cannot have volumes or aggregates on any internal drives. If you want to upgrade a FAS22xx system with volumes or aggregates on internal drives, you can upgrade by using the procedure *Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP 8.3 by moving volumes*.

Note: You can upgrade a FAS2220A system under one of the following circumstances:

- The system has no internal storage.

- The system has internal SATA drives or SSDs and you plan to transfer the drives to a disk shelf attached to the new system.

If you have a FAS2220A system with internal SAS drives, you can upgrade the system using the procedure *Upgrading the controller hardware on a pair of nodes running clustered Data ONTAP 8.3 by moving volumes*.

You can upgrade a FAS2240 system under the one of the following circumstances:

- The system has no internal storage.
- The system has internal storage and you plan to convert the FAS2240 to a disk shelf and attach it to the new system.

Note: If the new system has fewer slots than the original system, or if it has fewer or different types of ports, you might need to add an adapter to the new system. See the *Hardware Universe* on the NetApp Support Site for details about specific platforms.

Licensing in Data ONTAP 8.3

When you set up a cluster, the setup wizard prompts you to enter the cluster base license key. However, some features require additional licenses, which are issued as *packages* that include one or more features. Each node in the cluster must have its own key for each feature to be used in the cluster.

Starting with Data ONTAP 8.3, all license keys are 28 uppercase alphabetic characters in length. If you need to install a license apart from a Data ONTAP upgrade, you must enter the license key in the 28-character uppercase alphabetic format.

If you do not have new license keys, currently licensed features in the cluster will be available to the new controller. However, using features unlicensed on the controller might put you out of compliance with your license agreement, so you should install the new license key or keys on the new controller. You can obtain new license keys on the NetApp Support Site in the *My Support* section under *Software licenses*. If the site does not have the license keys you need, contact your NetApp sales representative.

For detailed information about licensing, see the *Clustered Data ONTAP System Administration Guide for Cluster Administrators* and the KB article *Data ONTAP 8.2 and 8.3 Licensing Overview and References* on the NetApp Support Site.

Storage Encryption

Storage Encryption is available beginning with clustered Data ONTAP 8.2.1. The original nodes or the new nodes might be enabled for Storage Encryption. In that case, you need to take additional steps in this procedure to ensure that Storage Encryption is set up properly.

If you want to use Storage Encryption, all the disk drives associated with the nodes must have self-encrypting disk drives.

Required tools and documentation

You must have specific tools to install the new hardware, and you need to reference other documents during the upgrade process. You also need to record information essential to completing the controller upgrade; a worksheet is provided to record information.

You need the following tools to perform the upgrade:

- Grounding strap
- #2 Phillips screwdriver

You need information in the following documents, which are available from the NetApp Support Site at mysupport.netapp.com:

Note: If you cannot access the NetApp Support Site when you are installing the new node, download the appropriate documents before beginning this procedure.

Document	Contents
<i>Clustered Data ONTAP Software Setup Guide</i>	Describes how to set up and configure NetApp systems
<i>Clustered Data ONTAP System Administration Guide for Cluster Administrators</i>	Contains instructions for obtaining information to set up the Service Processors (SPs) and performing other administrative tasks
<i>Clustered Data ONTAP Upgrade and Revert/Downgrade Guide</i>	Contains instructions for downloading and upgrading Data ONTAP
<i>Clustered Data ONTAP Physical Storage Management Guide</i>	Describes how to manage physical storage resources, using disks, RAID groups, and aggregates, and provides detailed information about Storage Encryption
<i>Clustered Data ONTAP Logical Storage Management Guide</i>	Describes how to efficiently manage your logical storage resources, using volumes, FlexClone volumes, files and LUNs, FlexCache volumes, deduplication, compression, qtrees, and quotas
<i>Installation and Setup Instructions for the model of the new nodes</i>	Contains instructions for installing and cabling the new system.
<i>Data ONTAP FlexArray Virtualization Installation Requirements and Reference Guide</i>	Contains cabling instructions and other information for FlexArray Virtualization systems
<i>Hardware Platform Monitoring Guide</i>	Contains information about system LEDs and system messages to aid troubleshooting
The appropriate disk shelf guide	Contains instructions for installing and monitoring disk shelves and replacing disk shelf devices
<i>Clustered Data ONTAP SAN Administration Guide</i>	Describes how to configure and manage iSCSI and FC protocols for SAN environments
<i>Clustered Data ONTAP SAN Configuration Guide</i>	Contains information about FC and iSCSI topologies and wiring schemes
<i>Clustered Data ONTAP High-Availability Configuration Guide</i>	Contains cabling instructions and other information for HA pairs
<i>MetroCluster Installation and Configuration Guide</i>	Describes how to install and configure the MetroCluster hardware and software components.

The NetApp Support Site also contains documentation about disk shelves, NICs, and other hardware that you might use with your system. It also contains the *Hardware Universe*, which provides information about the hardware that the new system supports.

Preparing to upgrade the node pair

Before you can upgrade the node pair, you need to ensure that the hardware is compatible and supported in clustered Data ONTAP 8.3, ensure that clustered Data ONTAP 8.3 is installed on both the old and new nodes, minimize the number of defined ports and failover rules, and capture system ID information.

Before you begin

If you are upgrading a pair of nodes in a switchless cluster, you must have converted them to a switched cluster (except for configurations where the new nodes are going to be FAS2520s) before performing the upgrade procedure. See the *Migrating from a switchless cluster to a switched Cisco®Nexus 5596, Nexus 5020, or Nexus 5010 cluster environment* or the *Migrating from a switchless cluster to a switched NetApp CN1610 cluster environment* on the NetApp Support Site at mysupport.netapp.com.

Steps

1. [Ensuring hardware and software compatibility between the original and new nodes](#) on page 5
2. [Checking license information and getting license keys](#) on page 5
3. [Rekeying disks for Storage Encryption](#) on page 5
4. [Recording original node port information](#) on page 7
5. [Planning port migration](#) on page 8
6. [Sending an AutoSupport message for the original nodes](#) on page 10
7. [Recording system ID information and shutting down the nodes](#) on page 10

Ensuring hardware and software compatibility between the original and new nodes

You need to ensure that the original hardware is compatible with the new hardware, and that both the original and new nodes are running Data ONTAP 8.3.

Steps

1. Check the *Hardware Universe* at hwu.netapp.com to verify that your existing and new hardware components are compatible and supported.

Note: You might need to add adapters to the new nodes.

2. Ensure that the original nodes have access to an HTTP server to download the latest software image from the NetApp Support Site.
3. Make sure that Data ONTAP 8.3 is installed on both the old nodes and the new nodes before performing the upgrade.

Both the old and new nodes must be running Data ONTAP 8.3 before the upgrade. Using a different version of the software on the new nodes results in storage subsystem panics. You can download Data ONTAP software images from the NOW site or request them from technical support. See the *Clustered Data ONTAP Upgrade and Revert/Downgrade Guide* for information about downloading and installing Data ONTAP.

Checking license information and getting license keys

If you want the new nodes to have the same licensed functionality as the original nodes, you need to see what licenses the original system has and then obtain new license keys for the new nodes.

Steps

1. Enter the following command to see which licenses are on the original system, and record the information:

```
system license show
```

You might need the information if you want the new system to have the same licensed features as the original system.

2. Obtain new license keys for the new nodes from the NetApp Support Site at mysupport.netapp.com.

Rekeying disks for Storage Encryption

If you used Storage Encryption on the original nodes and the new nodes have encryption-enabled disks, you must make sure that the original nodes' disks are correctly keyed.

About this task

Contact technical support to perform an optional step to preserve the security of the encrypted drives by rekeying all drives to a known authentication key.

Steps

1. On one of the nodes, access the nodeshell:

system node run -node node_name

The nodeshell is a special shell for commands that take effect only at the node level.

2. Display the status information to check for disk encryption:

disk encrypt show

Example

The system displays the key ID for each self-encrypting disk, as shown in the following example:

```
node> disk encrypt show
Disk      Key ID
0c.00.1   0x0
0c.00.0   080CF0C8000000000010000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.3   080CF0C8000000000010000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.4   080CF0C8000000000010000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.2   080CF0C8000000000010000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.5   080CF0C8000000000010000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
...
```

If you get the following error message, proceed to the “Preparing for netboot” section; if you do not get an error message, continue with these steps.

```
node> disk encrypt show
ERROR: The system is not configured to run this command.
```

3. Examine the output of the `disk encrypt show` command, and if any disks are associated with a non-MSID key, rekey them to an MSID key by taking one of the following actions:

- To rekey disks individually, enter the following command, once for each disk:

disk encrypt rekey 0x0 disk_name

- To rekey all disks at once enter the following command:

disk encrypt rekey 0x0 *

4. Verify that all the self-encrypting disks are associated with an MSID:

disk encrypt show

Example

The following example shows the output of the `disk encrypt show` command when all self-encrypting disks are associated with an MSID:

```
node> disk encrypt show
Disk      Key ID
-----
0b.10.23   0x0
0b.10.18   0x0
0b.10.0    0x0
0b.10.12   0x0
0b.10.3    0x0
0b.10.15   0x0
0a.00.1    0x0
0a.00.2    0x0
-----
Locked?
No
No
Yes
Yes
No
No
Yes
Yes
```

5. Obtain an IP address for the external key management server.

See the *Clustered Data ONTAP Software Setup Guide* for more information about the external key management server.

- Exit the nodeshell and return to the cluster shell:

```
exit
```

- Repeat Step 1 through Step 6 on the second node.

Recording original node port information

Before you shut down and retire the original nodes, record the information about their cluster network and management ports. You need that information later in the procedure when you map ports and LIFs from the original nodes to the new ones.

Steps

- Enter the following command on both node1 and node2 and record the information about the shelves, numbers of disks in each shelf, flash storage details, memory, NVRAM, and network cards from the output:

```
run -node node_name sysconfig
```

Note: You can use this information to identify parts or accessories that you might want to transfer to other nodes. If you do not know whether the nodes are V-Series systems or have FlexArray Virtualization, you can learn that from the output as well.

- If you are upgrading a V-Series system or a system with FlexArray Virtualization software, capture information about the topology of the original nodes by entering the following command and recording the output:

```
storage array config show -switch
```

Example

The system displays topology information, as show in the following example:

```
cluster:>> storage array config show -switch
```

Node	LUN Group	LUN Count	Array Name	Array Target Port	Target Side Switch Port	Initiator Side Switch Port	Initiator
node1	0	50	IBM_1818FAStT_1	205700a0b84772da	vgbr6510s164:5	vgbr6510s164:3	
				206700a0b84772da	vgbr6510s164:6	vgbr6510s164:4	2b
				207600a0b84772da	vgbr6510s163:6	vgbr6510s163:1	0c
node2	0	50	IBM_1818FAStT_1	205700a0b84772da	vgbr6510s164:5	vgbr6510s164:1	0d
				206700a0b84772da	vgbr6510s164:6	vgbr6510s164:2	2b
				207600a0b84772da	vgbr6510s163:6	vgbr6510s163:3	0c
				208600a0b84772da	vgbr6510s163:5	vgbr6510s163:4	2a

7 entries were displayed.

- Find the cluster network and node-management ports on the nodes by entering the following command on either controller:

```
network interface show -role cluster,node-mgmt
```

The system displays the LIFs for the node and other nodes in the cluster, as shown in the following example:

```
cluster:>> network interface show -role cluster,node-mgmt
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
node0	clus1	up/up	10.225.249.77/25	node0	e0a	true
	clus2	up/up	10.225.249.78/25	node0	e0b	true
	mgmt1	up/up	10.225.248.88/25	node0	e0M	true
node1	clus1	up/up	10.225.249.79/25	node1	e0a	true
	clus2	up/up	10.225.249.80/25	node1	e0b	true
	mgmt1	up/up	10.225.248.90/25	node1	e0M	true

6 entries were displayed.

- Capture or record the information in the output to use later in the procedure.
- Enter the following command on both nodes:

```
network port show -node node_name -type physical
```

Example

The system displays the physical ports on the node, as shown in the following example for node0:

```
cluster::> network port show -node node0 -type physical
```

Node	Port	IPspace	Broadcast Domain	Link	MTU	Speed (Mbps) Admin/Oper
node0	e0M	Default	10.225.248.88/24	up	1500	auto/100
	e0a	Default	-	up	1500	auto/1000
	e0b	Default	-	up	1500	auto/1000
	e1a	Cluster	Cluster	up	9000	auto/10000
	e1b	Cluster	Cluster	up	9000	auto/10000

5 entries were displayed.

- Record the ports and their broadcast domains.

The ports will need to be mapped to the correct broadcast domain later in the procedure.

Planning port migration

You need to identify the cluster, data, node-management, and physical ports on both the original and the new nodes. If necessary, you need to decide how you will migrate the ports on the original nodes to the ports on the new nodes later in the procedure.

About this task

Different platforms have different ports, so you might need to change the old nodes' port and LIF configuration to be compatible with the new nodes' configuration before you power down the old nodes. This is because the new nodes will replay the same configuration when they boot unless that configuration is changed before the upgrade.

For example, if the node-management LIF is hosted on port e1a of the old node, then booting the new node will cause it to also host the node management LIF on port e1a. However, if you want the new node to have its node management LIF on port e4a, then you must make the change on the original node before you boot it. If the two nodes have ports in common, you can choose the same port on the new node that is on the original node. If they do not, you need to choose the port on the old node that corresponds to the correct port on the new node.

Steps

- Identify the ports used for the data and management LIFs and record the information. You can modify these if you need to. You can also put data and node management LIFs in the same domain.

The default cluster ports on each platform are:

- FAS2220A: e1a, e1b
- FAS2240-2A: e1a, e1b
- FAS2240-4A: e1a, e1b
- FAS2520A: e0c, e0e
- FAS2552A, FAS2554A: e0e, e0f
- 32xx: e1a, e2a

- 62xx: e0c, e0e
- FAS8020: e0a, e0b
- FAS8040/FAS8060/FAS8080: e0a, e0b, e0c, e0d

The data ports on each platform are:

- FAS2220A: e0a, e0b, e0c, e0d
- FAS2240-2A: e0a, e0b, e0c, e0d
- FAS2240-4A: e0a, e0b, e0c, e0d
- FAS25xxA: e0a, e0b, e0c, e0d, e0e, e0f
- 32xx: e1b, e2b
- 62xx: e0d, e0f
- FAS8020A: e0c/0c, e0d/0d, e0e, e0f
- FAS8040/FAS8060/FAS8080: e0e/0e, e0f/0f, e0g/0g, e0h/0h, e0i, e0j, e0k, e0l

The default node-management ports on each platform are:

- FAS2240-4A: e0M
- FAS2220A: e0M
- FAS2240-2A: e0M
- FAS2240-4A: e0M
- FAS25xxA: e0M
- 32xxFAS32xxA: e0M
- 62xxFAS62xxA: e0M
- FAS80xxA: e0M

2. Determine which LIFs are present on which ports by entering the `network interface show` and the `network port show` commands and noting their output.
3. Consult the list in Step 1 or the *Hardware Universe* for port information on the new nodes.
4. Use the `network port modify` commands to adjust network broadcast domains, making the old node's configuration compatible with the new node's configuration.

The goal of this step is to obtain a compatible configuration for at least one port of each role. Compatible means that the configuration settings of the old node can be applied to the new node to enable connectivity when the new node boots.

See the *Clustered Data ONTAP Commands: Manual Page Reference* for more information about the `network port modify` command. Also see the lists in Step 1 for each platform's port assignments.

5. Complete the following substeps on both nodes to ensure that physical ports can be mapped correctly later in the procedure:
 - a. Enter the following command to see if there are failover groups on the node other than `clusterwide`:

```
network interface failover-groups show -vserver vservice_name
```

Failover groups are sets of network ports present on the system. Because upgrading the controller hardware can change the location of physical ports, failover groups can be inadvertently changed during the upgrade.

- b. If there are failover groups present other than `clusterwide`, record the failover group names and the ports that belong to the failover groups.
- c. Enter the following command to see if there are any VLANs configured on the node:

```
network port vlan show -node node_name
```

VLANs are configured over physical ports. If the physical ports change, then the VLANs need to be re-created later in the procedure.

Example

The system displays VLANs configured on the node, as shown in the following example:

```
cluster::> network port vlan show -node
                Network Network
Node   VLAN Name Port   VLAN ID MAC Address
-----
node0  e1b-70  e1b   70     00:15:17:76:7b:69
```

- d. If there are VLANs configured on the node, take note of each network port and VLAN ID pairing.
- e. If there are VLANs configured on interface groups (ifgrps), remove the VLANs from the interface groups by entering the following command:

```
network port vlan delete -node node_name -port ifgrp -vlan-id VLAN_ID
```

- f. Enter the following command and examine its output to see if there are any interface groups configured on the node:

```
network port ifgrp show -node node_name -instance
```

- g. If any interface groups are configured on the node, record the names of the interface groups and the ports assigned to them, and then delete the ports by entering the following command, once for each port:

```
network port ifgrp remove-port -node node_name -ifgrp ifgrp_name -port port_name
```

Sending an AutoSupport message for the original nodes

Before you perform the upgrade, you should send an AutoSupport message for the original nodes. You send a second AutoSupport message after the upgrade for the new nodes.

Step

1. Send an AutoSupport message to NetApp for the original nodes by entering the following command, once for each node:

```
system node autosupport invoke -node node_name -type all -message "Upgrading node_name from platform_original to platform_new"
```

Recording system ID information and shutting down the nodes

You must disable high availability or storage failover on the original nodes, record their system IDs, and destroy their mailboxes. You then need to halt the nodes and turn off their power.

About this task

Perform all the steps in this section on both of the original nodes, completing each step on both nodes before proceeding to the next step.

Steps

1. Take one of the following actions:

If you have a...	Then...
Two-node cluster	Disable high availability and storage failover by entering the following commands on the node that you are upgrading: <pre>cluster ha modify -configured false</pre> <pre>storage failover modify -node <i>node_name</i> -enabled false</pre> <i>node_name</i> should be one of the nodes that you are upgrading.
Cluster with more than two nodes	Disable storage failover by entering the following command on the node that you are upgrading: <pre>storage failover modify -node <i>node_name</i> -enabled false</pre> <i>node_name</i> should be one of the nodes that you are upgrading.

- Halt the node by entering the following command at the command prompt:

```
system node halt -node node_name
```

- Boot to Maintenance mode by entering the following command at the boot environment prompt:

```
boot_ontap maint
```

The Maintenance mode prompt appears as shown in the following example:

```
*>
```

Note: After you select 5 or enter the `boot_ontap maint` command, you might see the following message:

```
In a High Availability configuration, you MUST ensure that the
partner node is (and remains) down, or that takeover is manually
disabled on the partner node, because High Availability
software is not started or fully enabled in Maintenance mode.

FAILURE TO DO SO CAN RESULT IN YOUR FILESYSTEMS BEING DESTROYED

NOTE: It is okay to use 'show/status' sub-commands such as
'disk show or aggr status' in Maintenance mode while the partner is up
Continue with boot?
```

If the message appears, enter `y` when prompted.

- Display disk ownership information for the nodes that you are upgrading by entering the following command at the Maintenance mode prompt:

```
disk show -v
```

Example

```
*> disk show -v
Local System ID: 118049495
DISK      OWNER                POOL      SERIAL NUMBER                HOME
----      -
0a.33    node1 (118049495)        Pool0     3KS6BN970000973655KL        node1 (118049495)
0a.32    node1 (118049495)        Pool0     3KS6BCKD000097363ZHK        node1 (118049495)
0a.36    node1 (118049495)        Pool0     3KS6BL9H000097364W74        node1 (118049495)
0a.41    node1 (118049495)        Pool0     3KS6BNEX000097364BRE        node1 (118049495)
0c.43    node1 (118049495)        Pool0     V5YEB4TA                      node1 (118049495)
0c.41    node1 (118049495)        Pool0     V5YE72KA                      node1 (118049495)
0c.42    node1 (118049495)        Pool0     V5YD47BA                      node1 (118049495)
0c.24    node0 (118048821)        Pool0     V5YEK8KA                      node0 (118048821)
```

0c.17	node0 (118048821)	Pool0	V5YEGB9A	node0 (118048821)
0a.19	node0 (118048821)	Pool0	V5Y2K8HA	node0 (118048821)
0c.20	node0 (118048821)	Pool0	V5YEE01A	node0 (118048821)
0c.22	node0 (118048821)	Pool0	V5YETY0A	node0 (118048821)

- Record the System ID in the output of the `disk show -v` command.

You will enter the System ID as the value for `original_sys` in Step 5 of the section *Reassigning disks from the original nodes to the new nodes*.

- If the nodes are not new and unused, enter the following command at the Maintenance mode prompt on one of the original nodes:

```
mailbox destroy local
```

The console displays the following message:

```
Destroying mailboxes forces a node to create new empty mailboxes, which clears any
takeover state, removes all knowledge of out-of-date plexes and mirrored volumes, and
will prevent management services from going online in 2-node cluster HA configurations.
Are you sure you want to destroy the local mailboxes?
```

- Enter the following command at the prompt to confirm that you want to destroy the local mailboxes:

```
y
```

The system displays the following message:

```
.....Mailboxes destroyed
Takeover On Reboot option will be set to ON after the node boots.
This option is ON by default except on setups that have iSCSI or FCP license.
Use "storage failover modify -node <nodename> -onreboot false" to turn it OFF.
*>
```

- Exit Maintenance mode for both nodes that you are upgrading by entering the following command at the Maintenance mode prompt of each node:

```
halt
```

- Turn off the power to the original nodes and then unplug them from the power source.

Upgrading the node pair

Upgrading the node pair includes installing, powering up, and booting the new nodes, reassigning disks from the original nodes to the new ones, updating Data ONTAP from flash, and then mapping cluster and node-management ports from the original nodes to the new ones. You then perform several additional steps to complete the upgrade.

Steps

- [Installing the new nodes](#) on page 13
- [Booting and setting up the new nodes](#) on page 13
- [Transferring internal storage from the original nodes](#) on page 18
- [Reassigning disks from the original nodes to the new nodes](#) on page 20
- [Restoring configuration information from the root volume](#) on page 21
- [Mapping network ports](#) on page 22
- [Completing the upgrade](#) on page 25

Installing the new nodes

Installing the new nodes includes powering down and removing the original nodes and their chassis, installing the new nodes, and attaching power, console, and network connections to the new nodes.

Before you begin

You must have completed the steps in *Preparing to upgrade the node pair*.

Steps

1. Label and remove all cables from the original nodes.
2. Remove the chassis containing the original nodes.
3. If needed, install any adapters in the new nodes.

Follow the instructions in the appropriate adapter installation guide.

4. Install the new nodes, following the instructions in the *Installation and Setup Instructions* for the platform.

Note: Do not attach disk selves from the old nodes to the new nodes at this point. You perform this task in a later step.

Note: You need to perform Step 5 and Step 6 on both nodes.

5. Attach power and console connections to one of the new nodes.

For information about power and console connections, see the *Installation and Setup Instructions* for the node model.

6. Attach the network cables.

Booting and setting up the new nodes

You need to power on the new nodes, set the date, make sure that the nodes can run clustered Data ONTAP, boot Data ONTAP, and configure the new nodes. You then transfer any remaining cables from the original nodes to the new ones.

About this task

You need to perform the steps in this section on both nodes, completing each step on one node and then on the second node before proceeding to the next step.

Steps

1. Turn on the power to the first new node, and then immediately press Ctrl-C at the console terminal to access the boot environment prompt.

Note: If you are upgrading to a system with both nodes in the same chassis, the other node also reboots. However, you can disregard booting the other node until Step 10.

Attention: When you boot the new node, you might see the following message:

```
WARNING: The battery is unfit to retain data during a power
         outage. This is likely because the battery is
         discharged but could be due to other temporary
         conditions.
         When the battery is ready, the boot process will
         complete and services will be engaged.
         To override this delay, press 'c' followed by 'Enter'
```

2. If you see the warning message in Step 1, take the following actions:

- a. Check for any console messages that might indicate a problem other than a low NVRAM battery, and, if necessary, take any required corrective action.
- b. Allow the battery to charge and the boot process to complete.

Attention: Do not override the delay; failure to allow the battery to charge could result in a loss of data.

3. Take one of the following actions:

If your system...	Description
Has disks and no back-end storage	Skip Step 4 through Step 9 and go to Step 10.
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<ol style="list-style-type: none"> a. Go to either the “Configuring FC ports” or “Configuring CNA ports” section, as appropriate for your system. b. Return to this section and complete the remaining steps, beginning with Step 4. <p>Important: You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot Data ONTAP on the V-Series system or system with FlexArray Virtualization software.</p>

4. Add the FC initiator ports of the new node to the switch zones.

If your system has a tape SAN, then you need zoning for the initiators. See your storage array and zoning documentation for instructions.

5. Add the FC initiator ports to the storage array as new hosts, mapping the array LUNs to the new hosts.

See your storage array and zoning documentation for instructions.

6. Adjust the World Wide Port Name (WWPN) values in the host or volume groups associated with array LUNs on the storage array.

Installing a new controller module changes the WWPN and World Wide Node Name (WWNN) values associated with each onboard FC port.

7. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWNN values.

8. At the Maintenance mode prompt, halt the system:

```
halt
```

The system stops at the boot environment prompt.

9. Verify that the array LUNs are now visible on the new node:

```
sysconfig -v
```

The system displays all the array LUNs visible to each of the FC initiator ports. If the array LUNs are not visible, you will not be able to reassign disks from node1 to node3 later in this section.

10. Take one of the following actions:

If the system you are upgrading to is in a...	Then...
Dual-chassis configuration (with controllers in different chassis)	Go to Step 11.

If the system you are upgrading to is in a...	Then...
---	---------

Single-chassis configuration (with controllers in the same chassis)

- a. Switch the console cable from the first new node to the other new node.
 - b. Turn on the power to the second new node, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt. The power should already be on if both controllers are in the same chassis.
Note: Leave the second new node at the boot environment prompt; you return to this procedure and repeat these steps after the first new node is installed.
 - c. If you see the warning message displayed in Step 1, follow the instructions in Step 2.
 - d. Switch the console cable back from the first new node to the other new node.
 - e. Go to Step 11.
-

11. At the boot environment prompt, enter the following command:

```
set-defaults
```

12. At the boot environment prompt, enter the following command:

```
setenv bootarg.init.boot_clustered true
```

13. Configure the netboot connection by choosing one of the following actions.

Note: You should use the management port and IP as the netboot connection. Do not use a data LIF IP address, or a data outage might occur while the upgrade is being performed.

If DHCP is...	Then...
---------------	---------

Running

Configure the connection automatically by entering the following command at the boot environment prompt:

```
ifconfig e0M -auto
```

In a single-chassis configuration (with controllers in the same chassis)

Manually configure the connection by entering the following command at the boot environment prompt:

```
ifconfig e0M -addr=filer_addr mask=netmask -gw=gateway -  
dns=dns_addr -domain=dns_domain
```

filer_addr is the IP address of the storage system.

netmask is the network mask of the storage system.

gateway is the gateway for the storage system.

dns_addr is the IP address of a name server on your network.

dns_domain is the Domain Name Service (DNS) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name.

Note: Other parameters might be necessary for your interface. Entering the `help ifconfig` command at the firmware prompt provides details.

14. Perform netboot by entering the following command on the first new node:

```
netboot http://path_to_the_web-accessible_directory/netboot/kernel
```

The *path_to_the_web-accessible_directory* should lead to where you downloaded the `netboot.tgz` file in Step 1a in the “Preparing for netboot” section.

Note: Do not interrupt the boot.

- From the boot menu, select option (7) **Install new software first**.

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

Note: Disregard the following message: "This procedure is not supported for Non-Disruptive Upgrade on an HA pair". The note applies to nondisruptive upgrades of Data ONTAP, not to upgrades of controllers.

- If you are prompted to continue the procedure, enter **y**, and when prompted for the package, enter the URL `http://path_to_the_web-accessible_directory/image.tgz`.
- Complete the following substeps:

- Enter **n** to skip the backup recovery when you see the following prompt:

```
Do you want to restore the backup configuration now? {y|n}
n
```

- Reboot by entering **y** when you see the following prompt:

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

The controller module reboots but stops at the boot menu because the boot device was reformatted and the configuration data needs to be restored.

- Boot the new node:

```
boot_ontap
```

- Interrupt the boot by pressing Ctrl-C.

The system displays the boot menu.

- Select (5) **Maintenance mode boot** by entering **5**, and then enter **y** when prompted to continue with the boot.

- Before continuing to Step 22, go to the sections "Configuring FC ports" or "Configuring CNA ports" to make any necessary changes to the FC or CNA ports on the node.

Make the changes recommended in those sections, reboot the node, and go into Maintenance mode.

- Find the system ID of node3:

```
disk show -a
```

Example

The system displays the system ID of the node and information about its disks, as shown in the following example:

```
*> disk show -a
Local System ID: 536881109
DISK      OWNER                                POOL  SERIAL NUMBER HOME                                DR HOME
-----
0b.02.23 nst-fas2520-2(536880939) Pool10 KPG2RK6F      nst-fas2520-2(536880939)
0b.02.13 nst-fas2520-2(536880939) Pool10 KPG3DE4F      nst-fas2520-2(536880939)
0b.01.13 nst-fas2520-2(536880939) Pool10 PPG4KLAA      nst-fas2520-2(536880939)
.....
0a.00.0   (536881109) Pool10 YFKSX6JG      (536881109)
.....
```

- Reassign the new node's spares, disks belonging to the root, and any non-root aggregates that were not relocated to the node earlier:

```
disk reassign -s new_node_sysid -d node3_sysid -p node2_sysid
```


When you run the `disk reassign` command on the new node, the `-p` option is the `node2_sysid`; when you run the `disk reassign` command on the other node, the `-p` option is the `node3_sysid`.

For the `new _node_sysid` value, use the information captured in Step 2 of the “Recording original node port information” section. To obtain the value for `node3_sysid`, use the `sysconfig` command.

The `disk reassign` command reassigns only those disks for which `node1_sysid` is the current owner.

The system displays the following message:

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

24. Enter `n`.

The system displays the following message:

```
After the node becomes operational, you must perform a takeover and giveback of the HA
partner node to ensure disk reassignment is successful.
Do you want to continue (y/n)? y
```

25. Enter `y`.

The system displays the following message:

```
Disk ownership will be updated on all disks previously belonging to Filer with sysid
<sysid>.
Do you want to continue (y/n)? y
```

26. Enter `y`.

27. Verify that the controller and chassis are configured as `ha`:

```
ha-config show
```

Example

The following example shows the output of the `ha-config show` command:

```
*> ha-config show
  Chassis HA configuration: ha
  Controller HA configuration: ha
```

Systems record in a PROM whether they are in an HA pair or stand-alone configuration. The state must be the same on all components within the stand-alone system or HA pair.

If the controller and chassis are not configured as `ha`, use the following commands to correct the configuration: `ha-config modify controller ha` and `ha-config modify chassis ha`.

If you have a MetroCluster configuration, use the following commands to modify the controller and chassis: `ha-config modify controller mcc` and `ha-config modify chassis mcc`.

28. Destroy the mailboxes on `node3`:

```
mailbox destroy local
```

The console displays the following message:

```
Destroying mailboxes forces a node to create new empty mailboxes,
which clears any takeover state, removes all knowledge of
out-of-date plexes of mirrored volumes, and will prevent
management services from going online in 2-node cluster HA
configurations. Are you sure you want to destroy the local
mailboxes?
```

29. Enter **y** at the prompt to confirm that you want to destroy the local mailboxes.

30. Exit Maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

31. On the other node, check the system date, time, and time zone:

```
date
```

32. At the boot environment prompt, check the system date:

```
show date
```

33. If the date is incorrect, set the date by entering the following command at the boot environment prompt:

```
set date mm/dd/yyyy
```

34. At the boot environment prompt, verify that the *bootarg* variable was set correctly:

```
printenv
```

35. Transfer all remaining cables from the old node to the corresponding ports on the new node.

This includes Fibre Channel or other external disk shelf cables, and Ethernet cables. You need to ensure that all the cables are correctly connected. For information about cabling, see the installation and setup instructions for the node model.

Transferring internal storage from the original nodes

You can transfer internal SATA drives or SSDs from a FAS2220A system to a disk shelf attached to the new nodes or convert a FAS2240 system to a disk shelf and attach it to the new nodes.

Before you begin

You must have completed the previous steps in this procedure up to this point.

Choices

- [Moving internal disk drives from a FAS2220 system](#) on page 18
- [Converting the FAS2240 system to a disk shelf and attaching it to the new system](#) on page 20

Moving internal disk drives from a FAS2220 system

After you have installed and set up the new node, you can move the internal disk drives from a FAS2220 system with SATA drives or SSDs to a disk shelf attached to the new system.

Before you begin

You must have done the following before proceeding with this section:

- Made sure that the SATA or SSD drive carriers from the FAS2220 system are compatible with the new disk shelf
Check the *Hardware Universe* on the NetApp Support Site for compatible disk shelves
- Made sure that there is a compatible disk shelf attached to the new system

- Made sure that the disk shelf has enough free bays to accommodate the SATA or SSD drive carriers from the FAS2220 system


About this task

You cannot transfer SAS disk drives from a FAS2220 system to a disk shelf attached to the new nodes.

Steps

1. Gently remove the bezel from the front of the system.
2. Press the release button on the left side of the drive carrier.

The following illustration shows a disk drive with the release button located on the left of the carrier face:




The cam handle on the carrier springs open partially, and the carrier releases from the midplane.

3. Pull the cam handle to its fully open position to unseat the carrier from the midplane and gently slide the carrier out of the disk shelf.

Attention: Always use two hands when removing, installing, or carrying a disk drive. However, do not place your hands on the disk drive boards exposed on the underside of the carrier.

The following illustration shows a carrier with the cam handle in its fully open position:



4. With the cam handle in the open position, insert the carrier into a slot in the new disk shelf, firmly pushing until the carrier stops.

Caution: Use two hands when inserting the carrier.

5. Close the cam handle so that the carrier is fully seated in the midplane and the handle clicks into place.
Be sure you close the handle slowly so that it aligns correctly with the face of the carrier.
6. Repeat Step 2 through Step 5 for all of the disk drives that you are moving to the new system.

Converting the FAS2240 system to a disk shelf and attaching it to the new system

After you complete the upgrade, you can convert the FAS2240 system to a disk shelf and attach it to the new system to provide additional storage.

Before you begin

You must have upgraded the FAS2240 system before converting it to a disk shelf. The FAS2240 system must be powered down and uncabled.

Steps

1. Replace the controller modules in the FAS2240 system with IOM6 modules.
2. Set the disk shelf ID.
Each disk shelf, including the FAS2240 chassis, requires a unique ID.
3. Reset other disk shelf IDs as needed.
4. Turn off power to any disk shelves connected to the new nodes, and then turn off power to the new nodes.
5. Cable the converted FAS2240 disk shelf to a SAS port on the new system, and, if you are using ACP cabling, to the ACP port on the new system.
Note: If the new system does not have a dedicated onboard network interface for ACP for each controller, you must dedicate one for each controller at system setup. See the *Installation and Setup Instructions* for the new system and the *Universal SAS and ACP Cabling Guide* for cabling information. Also consult the *Clustered Data ONTAP High-Availability Configuration Guide*.
6. Turn on power to the converted FAS2240 disk shelf and any other disk shelves attached to the new nodes.
7. Turn on the power to the new nodes and then interrupt the boot process on each node by pressing Ctrl-C to access the boot environment prompt.

Reassigning disks from the original nodes to the new nodes

You must reassign the disks that belong to the original nodes to the new nodes.

About this task

You must perform the steps in this section on both nodes, completing each step on one node and then the other node before going on to the next step.

Steps

1. Boot Data ONTAP on the new node by entering the following command at the boot environment prompt:
`boot_primary maint`
2. If the new system has a root aggregate, take the root aggregate of the new system offline and destroy it by entering the following commands:
`aggr offline aggr0`
`aggr destroy aggr0`
Note: The name of the root aggregate might not be `aggr0` so make the appropriate changes as required to the previous commands.
3. On the new node, display the new node system ID by entering the following command at the Maintenance mode prompt:
`disk show -v`

Example

```
*> disk show -v
Local System ID: 101268854
...
```

- Record the new node system ID.
- Reassign the node's spares, disks belonging to root, and any SFO aggregates by entering the following command:

```
disk reassign -s original_sysid -d new_sysid -p partner_sysID
```

The `disk reassign` command will only reassign those disks for which `original_sysid` is the current owner.

The system displays a message and prompts as shown in the following example:

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

After the node becomes operational, you must perform a takeover and giveback of the HA
partner node to ensure disk reassignment is successful.
Do you want to continue (y/n)?y
```

- Enter `y` when prompted.

The system displays the following message:

```
The system displays the following message:
Disk ownership will be updated on all disks previously belonging to Filer with sysid
<sysid>.
Do you want to continue (y/n)? y
```

- Enter `y` when prompted.
- Make sure that the old system's root aggregate is set as Options `root` by entering the following command:

```
aggr status
```

You should see output similar to the following:

```
*> aggr status
      Aggr State           Status           Options
      aggr0 online        raid_dp, aggr   root
                        64-bit
```

Besides checking for the appropriate root aggregate, make sure that any other aggregates are online as well.

- Enter the following command at the Maintenance mode prompt:

```
halt
```

Restoring configuration information from the root volume

After you install and boot the new nodes and reassign disks, you must restore configuration information from the root volume to the boot devices.

About this task

You need to perform the steps in this section on both nodes, completing each step on one node and then the other before going on to the next step.

Steps

1. Enter the following command at the boot environment prompt:

```
boot_primary menu
```

The node reboots and then restarts and displays the boot menu.

2. From the boot menu, select **(6) Update flash from backup config.**, as shown in the following example:

```
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.
Selection (1-8)? 6
```

The system displays the following message:

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

3. Enter **y**.

The update flash process runs for several minutes, and then the system reboots. The startup process then asks you to confirm the system ID mismatch.

4. Confirm the mismatch as shown in the following example:

```
WARNING: System id mismatch. This usually occurs when replacing CF or NVRAM cards!
Override system id? {y|n} [n] y
```

The startup sequence proceeds normally.

Mapping network ports

You need to make sure that the physical ports on the original node map correctly to the physical ports on the new node, which allow the new node to communicate with other nodes in the cluster and with the network after the upgrade.

Before you begin

You can capture information about the ports for the new node from the *Hardware Universe* at hwu.netapp.com. You use the information later in this section.

About this task

The software configuration of the new node must match the physical connectivity of the node, and IP connectivity must be restored before you continue with the upgrade.

Port settings might vary, depending on the model of the nodes. You must make the original node's port and LIF configuration compatible with what you plan the new node's configuration to be. This is because the new node replays the same configuration when it boots, which means that when you boot the new node Data ONTAP tries to host LIFs on the same ports that were used on the original node.

If the physical ports on the original node do not map directly to the physical ports on the new node, software configuration changes will be required to restore cluster, management, and network connectivity after the boot. In addition, if the cluster ports

on the original node do not directly map to the cluster ports on the new node, the new node may not automatically rejoin quorum when it is rebooted until a software configuration change is made to host the cluster LIFs on the correct physical ports.

Steps

1. Record all of the original node's cabling information for that node, the ports, broadcast domains, and IPspaces, in this table:

LIF	Original node ports	Original node IPspaces	Original broadcast domains	New node ports	New node IPspaces	New node broadcast domains	SAN/NAS
Cluster 1							NA
Cluster 2							NA
Cluster 3							NA
Cluster 4							NA
Cluster 5							NA
Cluster 6							NA
Node management							NA
Cluster management							NA
Data 1							
Data 2							
Data 3							
Data 4							

See the “Recording original node port information” section for the steps to obtain this information.

2. Record all the cabling information for the new node, the ports, broadcast domains, and IPspaces in the previous table using the same procedure; see the “Recording original node port information” section for the steps to obtain this information.
3. Place the new node into quorum, perform the following steps:
 - a. Boot the new node. See the “Booting and setting up the new nodes” section to boot the node if you have not already done so.

- b. Add the correct ports to the Cluster broadcast domain:

```
network port modify -ipspace Cluster -mtu 9000
```

Example

This example adds Cluster port e1b on node-new:

```
network port modify -node node-new -port e1b -ipspace Cluster -mtu 9000
```

- c. Migrate the LIFs to the new ports, once for each LIF:

```
network interface migrate -vserver vs_server_name -lif lif_name -source-node node-new -destination-node node-new -destination-port port_name
```

Note: If you cannot migrate a cluster LIF, make sure the node has another network port that can communicate on the Cluster network and migrate the cluster LIF to that network port. After the node has joined quorum, you try a reboot to resolve the issue of Data ONTAP not having discovered the network port.

If you cannot migrate a data LIF, delete the LIF and re-create it on a functioning network port.

Note: You cannot migrate SAN data LIFs. Those LIFs need to be taken offline before you can change their home ports.

- d. Modify the home port of the cluster LIFs:

```
network interface modify -vserver Cluster -lif lif_name -home-port port_name
```

- e. Remove the old ports from the Cluster broadcast domain:

```
network port broadcast-domain remove-ports
```

Example

This command removes port e0d on node-new:

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast-domain Cluster -ports node-new:e0d
```

- f. Verify that node-new has rejoined quorum:

```
cluster show -node node-new -fields health
```

4. If there were any ports on the original node that no longer exist on the new node, delete them:

```
network port delete -node node_name -port port_name
```

5. Adjust the node-management broadcast domain and migrate the node-management and cluster-management LIFs, if necessary, as follows:

- a. Display the home port of a LIF:

```
network interface show -fields home-node,home-port
```

- b. Display the broadcast domain containing this port:

```
network port broadcast-domain show -ports node_name:port_name
```

- c. Add or remove ports from broadcast domains:

```
network port broadcast-domain add-ports
```

```
network port broadcast-domain remove-ports
```

- d. To modify a LIF's home port:

```
network interface modify -vserver vserver -lif lif_name -home-port port_name
```

6. Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary, using the same commands shown in Step 5.

7. Adjust any other broadcast domains and migrate the data LIFs, if necessary, using the same commands shown in Step 5.

8. Adjust all the LIF failover groups:

```
network interface modify -failover-group failover_group -failover-policy failover_policy
```

Example

The following command sets the failover policy to `broadcast-domain-wide` and uses the ports in failover group `fg1` as failover targets for LIF `data1` on node3:

```
network interface modify -vserver node3 -lif data1 failover-policy broadcast-domain-wide -failover-group fg1
```

See the “Configuring failover settings on a LIF” section in the *Clustered Data ONTAP Network Management Guide* or *Clustered Data ONTAP Commands: Manual Page Reference* for more information.

9. Verify the changes on the new node:

```
network port show -node node-new
```

Completing the upgrade

To complete the upgrade, you need to delete any unused ports and LIFs from the new nodes, re-enable storage failover or high availability, rename the nodes if necessary, and configure the SPs.

About this task

In this section, complete Step 1 through Step 8 on both of the new nodes.

Steps

1. At the command prompt, enter:

```
network interface show
```

2. Examine the output of the `network interface show` command to verify that the logical interfaces are operational:

Example

```
cluster::> network interface show
Server Logical      Status      Network      Current Current  Is
----- Interface    Admin/Oper  Address/Mask Node         Port      Home
-----
node1
  clus1      up/up      10.90.200.100/21  node1      e0c      true
  clus2      up/up      10.90.200.101/21  node1      e0e      true
  mgmt1      up/down    10.98.201.97/21   node1      e0M      true
node2
  clus1      up/up      10.90.200.102/21  node2      e0c      true
  clus2      up/up      10.90.200.103/21  node2      e0e      true
  mgmt1      up/down    10.90.201.99/21   node2      e0M      true
vs9
  node1_data1 up/up      10.90.225.50/20   node1      e0d      true
  node1_data2 up/up      10.90.225.51/20   node1      e0f      true
  node2_data1 up/up      10.90.225.53/20   node2      e0d      true
  node2_data2 up/up      10.90.225.52/20   node2      e0f      true
10 entries were displayed.
```

Note: After you enter the `network interface show` command, you might not see the output because the node is completing startup. If you do not see the output, wait 5 to 10 seconds and then run the `network interface show` command again.

3. Delete any unused ports on the new nodes by completing the following substeps:
 - a. Access the advanced privilege level by entering the following command on either node:

```
set -privilege advanced
```

The system displays the following message:

```
Warning: These advanced commands are potentially dangerous; use them only when
directed to do so by NetApp personnel.
do you wish to continue? {y | n}:
```

- b. Enter `y`.
- c. Enter the following command, once for each port that you want to delete:

```
network port delete -node node_name -port port_name
```

d. Return to the admin level by entering the following command:

```
set -privilege admin
```

4. Delete any unused network interfaces on the new nodes by completing the following substeps:

If a LIF you want to remove is in a port set, you must delete the LIF from the port set before you can remove it.

a. If you are in a SAN environment, check to see if any unused LIFs are in port sets by entering the following command and examining its output:

```
lun portset show
```

b. If any unused LIFs are in port sets, remove the LIFs from the port sets by entering the following command:

```
lun portset remove -vserver vserver_name -portset portset_name -port-name port_name
```

c. Remove each unused LIF by entering the following command, once for each LIF:

```
network interface delete -vserver vserver_name -lif LIF_name
```

5. Re-enable storage failover or high availability on the pair of new nodes as needed by taking one of the following actions:

If you have a...	Then...
Two-node cluster	Re-enable high availability by entering the following command on one of the nodes: cluster ha modify -configured true
A cluster with more than two nodes	Re-enable storage failover by entering the following command on one of the nodes: storage failover modify -node node_name -enabled true

6. If the name of the nodes includes the platform name (for example, node-8080-1 and node-8060-2), rename the nodes and ensure that the new node names are reflected on both partners in the pair by entering the following command at the command prompt for both nodes:

```
system node rename -node current_name -newname new_name
```

See the *Clustered Data ONTAP System Administration Guide for Cluster Administrators* for more information about the `system node rename` command.

7. If the node names are changing, enter the following command to find the reporting nodes list:

```
lun mapping show -vserver vserver_name -path lun-path igroup igroup_name
```

Example

The system displays input similar to this example:

```

::*>lun mapping show -vserver node1 -path /vol/src_vol/src_lun_1_5 igroup san_ig
Vserver Name: node1
LUN Path: /vol/src_vol/src_lun_1_5
Volume Name: src_vol
Qtree Name: ""
LUN Name: src_lun_1_5
Igroup Name: san_ig
Igroup OS Type: windows
Igroup Protocol Type: iscsi
LUN ID: 5
Portset Binding Igroup: -
ALUA: true
Initiators: iqn.1991-05.com.microsoft:netappn-o6ntfii
LUN Node: ssan-8020-09a
Reporting Nodes: ssan-8020-09a, ssan-8020-09b <-----

```

```
Vserver UUID: b289ac6f-37b7-11e4-987c-00a09859c899
LUN UUID: fe52b4c3-19d1-4a52-912b-139816f5591e
Igroup UUID: df4cf0c7-37b7-11e4-987c-00a09859c899
Portset UUID: -
```

- To modify the reporting list to add the local HA pairs for each LUN:

```
lun mapping add-reporting-nodes -path * -local-nodes
```

- To remove old nodes from the reporting nodes list:

```
lun mapping remove-reporting-nodes -vserver vserver_name -path path_name -igroup igroup_name
```

- If you renamed the nodes in Step 6, check that the name change took effect on both nodes by entering the following command and examining its output:

```
system node show
```

The `system node show` command displays information about all the nodes in the cluster, as shown in the following example:

```
cluster::> system node show
Node   Health Eligibility Uptime           Model   Owner   Location
-----
node0  true   true             89 days 23:47 MODELXX  IT      Data Center 2
node1  true   true             15 days 22:37 MODELXX  Data Center 2
node2  true   true             15 days 23:00 MODELXX  Data Center 2
node3  true   true             15 days 22:37 MODELXX  Data Center 2
4 entries were displayed.
```

- If the name of the Storage Virtual Machine (SVM, formerly known as Vserver) includes a platform name, rename it by entering the following command:

```
vserver rename -vserver current_vserver_name -newname new_vserver_name
```

- Configure the SP by performing the following command on both nodes:

```
system service-processor network modify
```

See the *Clustered Data ONTAP System Administration Guide for Cluster Administrators* for information about the SPs and the *Clustered Data ONTAP Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

- Install new licenses on the new nodes by entering the following command on each node:

```
system license add license_code,license_code,license_code...
```

The parameter accepts a list of 28 upper-case alphabetic character keys. You can add one license at a time or multiple licenses, separating each license key with a comma or space.

- If you want to set up a switchless cluster on the new nodes, follow the instructions in *Migrating to a two-node switched cluster with Cisco® cluster switches* or *Migrating to a two-node switched cluster with NetApp CN1610 cluster switches* on the NetApp Support Site.

Performing post-upgrade tasks

After you complete the upgrade, you should set up AutoSupport and send a post-upgrade message to NetApp. You should also set up Storage Encryption on the new nodes if they are encryption-enabled. In addition, if one or both of the nodes has a unified target adapter, you need to configure adapter ports and might need to change the personality of the adapter.

Steps

1. [Setting up AutoSupport and sending a post-upgrade AutoSupport message](#) on page 28
2. [Setting up Storage Encryption on the new nodes](#) on page 28
3. [Configuring FC ports](#) on page 29
4. [Configuring CNA ports](#) on page 31
5. [Decommissioning the old system](#) on page 34

Setting up AutoSupport and sending a post-upgrade AutoSupport message

Once the upgrade is complete, you should set up AutoSupport on the new nodes and then send a post-upgrade AutoSupport message to NetApp.

Steps

1. Set up AutoSupport, following the instructions in the *Clustered Data ONTAP System Administration Guide for Cluster Administrators*.
2. Send a post-upgrade AutoSupport message to NetApp by entering the following command, once for each node:

```
system node autosupport invoke -node node_name -type all -message "node_name successfully upgraded from platform_old to platform_new"
```

Setting up Storage Encryption on the new nodes

If the new nodes have Storage Encryption enabled, you might need to complete a series of additional steps to ensure uninterrupted Storage Encryption functionality. These steps include collecting network information, obtaining private and public SSL certificates, and running the Storage Encryption setup wizard.

Before you begin

All the disks on the storage system must be encryption-enabled before you set up Storage Encryption on the new nodes.

About this task

You can skip this section if the system that you upgraded to does not have Storage Encryption enabled.

If you used Storage Encryption on the original system and migrated the disk shelves to the new system, you can reuse the SSL certificates that are stored on migrated disk drives for Storage Encryption functionality on the upgraded system. However, you should check that the SSL certificates are present on the migrated disk drives. If they are not present you will need to obtain them.

Note: Step 1 through Step 3 are only the overall tasks required for configuring Storage Encryption. You need to follow the detailed instructions for each task in the *Clustered Data ONTAP Software Setup Guide*.

Steps

1. Obtain and install private and public SSL certificates for the storage system and a private SSL certificate for each key management server that you plan to use.

Requirements for obtaining the certificates and instructions for installing them are contained in the *Clustered Data ONTAP Software Setup Guide*.

2. Collect the information required to configure Storage Encryption on the new nodes.

This includes the network interface name, the network interface IP address, and the IP address for external key management server. The required information is contained in the *Clustered Data ONTAP Software Setup Guide*.

3. Launch and run the Storage Encryption setup wizard, responding to the prompts as appropriate.
4. If you have not done so, repeat Step 1 through Step 3 on the other new node.

After you finish

See the *Clustered Data ONTAP Physical Storage Management Guide* for information about managing Storage Encryption on the updated system.

Configuring FC ports

If the new node has Fibre Channel (FC) ports—onboard or on an FC adapter—you must set port configurations on the node before you bring it into service because the ports are not preconfigured. If the ports are not configured, you might experience a disruption in service.

Before you begin

You must have the values of the FC port settings for the new node that you saved in the “Preparing to upgrade the node pair” and sections.

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard CNA ports or a CNA adapter, you configure them in the next section.

Important: If your system has storage disks, you enter the commands in this section at the cluster prompt. If you have a V-Series system connected to storage arrays, you enter the commands in this section in Maintenance mode.

Steps

1. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Go to Step 3.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 2.

2. If you have not already done so, halt the node:

```
system node halt -node node_name
```

3. Boot the new node and access Maintenance mode:

```
boot_ontap maint
```

4. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<code>system node hardware unified-connect show</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>ucadmin show</code>

The system displays information about all FC and converged network adapters on the system.

5. Compare the FC settings on the new node with the settings that you recorded from the original node.

6. Take one of the following actions:

If the default FC settings on the new nodes are...	Then...
---	----------------

The same as the ones you that captured on the original node	Go to Step 11.
---	----------------

Different from the ones that you captured on the original node	Go to Step 7.
--	---------------

7. Take one of the following actions:

If the system that you are upgrading...	Then...
--	----------------

Has storage disks	Modify the FC ports on the new node as needed: <ul style="list-style-type: none">To program target ports:<pre>system node hardware unified-connect modify -type -t target -adapter port_name</pre>To program initiator ports:<pre>system node hardware unified-connect modify -type -t initiator -adapter port_name</pre> <p>-type is the FC4 type: target or initiator.</p>
-------------------	---

Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Modify the FC ports on the new node as needed: <pre>ucadmin modify -m fc -t initiator -f adapter_port_name</pre> <p>-t is the FC4 type, target or initiator.</p> <p>Note: The FC ports must be programmed as initiators.</p>
--	---

8. Take one of the following actions:

If the system that you are upgrading...	Then...
--	----------------

Has storage disks	Verify the new settings: <pre>system node hardware unified-connect show</pre>
-------------------	---

Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Verify the new settings: <pre>ucadmin show</pre>
--	--

9. Exit Maintenance mode:

```
halt
```

10. After you enter the command, wait until the system stops at the boot environment prompt.

11. Boot the new node and access Maintenance mode by entering the following command at the boot environment prompt:

```
boot_ontap maint
```

12. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<ul style="list-style-type: none"> • If the new node has a CNA card or CNA onboard ports, go to the “Configuring CNA ports” section. • If the new node does not have a CNA card or CNA onboard ports, skip the “Configuring CNA ports” section, and go to the “Mapping network ports” section.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ul style="list-style-type: none"> • If the new node has a CNA card or CNA onboard ports, go to the “Configuring CNA ports” section. • If the new node does not have a CNA card or CNA onboard ports, skip the “Configuring CNA ports” section, and return to Step 4 in the “Booting and setting up the new node” section.

Configuring CNA ports

If a node has onboard CNA ports or a CNA card, you must check the configuration of the ports and possibly reconfigure them, depending on how you want to use the upgraded system.

Before you begin

You must have the correct SFP+ modules for the CNA ports.

About this task

CNA ports can be configured into native Fibre Channel (FC) mode or CNA mode. FC mode supports FC initiator and FC target; CNA mode allows concurrent NIC and FCoE traffic over the same 10-GbE SFP+ interface and supports FC target.

Note: NetApp marketing materials might use the term *UTA2* to refer to CNA adapters and ports. However, the CLI and product documentation use the term *CNA*.

CNA ports might be on an adapter or onboard the controller and have the following configurations:

- CNA cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- CNA cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard CNA ports on new controllers are configured before shipment to have the personality you request.

However, you should check the configuration of the CNA ports on the node and change them, if necessary.

Steps

1. If you have not already done so, halt the node:

```
system node halt -node node_name
```

2. Access Maintenance mode:

```
boot_ontap maint
```

3. Check how the ports are currently configured by entering one of the following commands on one of the new nodes:

If the system you are upgrading...	Then...
Has storage disks	Enter the following command: <code>system node hardware unified-connect show</code>

If the system you are upgrading...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Enter the following command: ucadmin show

The system displays output similar to the following examples:

```
cluster1::> system node hardware unified-connect show
Node   Adapter  Current Mode   Current Type   Pending Mode   Pending Type   Admin Status
-----
f-a    0e       fc     initiator -      -      online
f-a    0f       fc     initiator -      -      online
f-a    0g       cna    target  -      -      online
f-a    0h       cna    target  -      -      online
f-b    0e       fc     initiator -      -      online
f-b    0f       fc     initiator -      -      online
f-b    0g       cna    target  -      -      online
f-b    0h       cna    target  -      -      online
8 entries were displayed.
```

```
node*> ucadmin show
Adapter  Current Mode   Current Type   Pending Mode   Pending Type   Status
-----
0e       fc     initiator -      -      online
0f       fc     initiator -      -      online
0g       cna    target  -      -      online
0h       cna    target  -      -      online
0e       fc     initiator -      -      online
0f       fc     initiator -      -      online
0g       cna    target  -      -      online
0h       cna    target  -      -      online
*>
```

- If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module.
- Examine the output of the `ucadmin show` or `system node hardware unified-connect show` command and determine whether the CNA ports have the personality you want.
- Take one of the following actions:

If the CNA ports ...	Then...
Do not have the personality that you want	Go to Step 7.
Have the personality that you want	Go to Step 9.

- If the CNA adapter is online, take it offline by entering one of the following commands:

If the system that you are upgrading...	Then...
Has storage disks	<ul style="list-style-type: none"> If the adapter is in initiator mode, enter the following command: system node run -node node-name disable adapter adapter-name If the adapter is in target mode, enter the following command: fcp adapter modify -node node-name -adapter adapter-name -state down

If the system that you are upgrading...	Then...
Is a V-Series system or has FlexArray Software and is attached to storage arrays	Enter the following command: storage disable adapter <i>adapter-name</i>

Adapters in target mode are automatically placed offline in Maintenance mode.

8. If the current configuration does not match the desired use, enter the following commands to change the configuration as needed:

If the system that you are upgrading...	Then...
Has storage disks	Enter the following command: system node hardware unified-connect modify -node <i>node-name</i> -adapter <i>adapter-name</i> -mode fc cna -type target initiator
Is a V-Series system or has FlexArray Software and is attached to storage arrays	Enter the following command: ucadmin modify -m fc cna -t initiator target <i>adapter-name</i>

In either command:

- -m or -mode is the personality mode, fc or 10GbE cna.
- -t or -type is the FC4 type, target or initiator.

Note: You need to use an FC initiator for tape drives, FlexArray Virtualization systems, and Fabric MetroCluster. You also need to use FC initiator for stretch MetroCluster if you are using a FibreBridge6500N bridge. You need to use an FC target for SAN clients.

9. Verify the settings by entering one of the following commands and examining its output:

If the system that you are upgrading...	Then...
Has storage disks	Enter the following command: system node hardware unified-connect show
Is a V-Series system or has FlexArray Software and is attached to storage arrays	Enter the following command: ucadmin show

Example

The output in the following examples show that the FC4 type of adapter 1b is changing to initiator and that the mode of adapters 2a and 2b is changing to cna:

```
cluster1::> system node hardware unified-connect show
Node  Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Status
-----
f-a   1a       fc      initiator  -      -      online
f-a   1b       fc      target    -      initiator  online
f-a   2a       fc      target    cna     -      online
f-a   2b       fc      target    cna     -      online
4 entries were displayed.
```

```
node> ucadmin show
Adapter  Current Mode  Current Type  Pending Mode  Pending Type  Status
-----
```

```

1a      fc      initiator -      -      online
1b      fc      target   -      initiator online
2a      fc      target   cna    -      online
2b      fc      target   cna    -      online
node>

```

10. Enter the following command:

```
halt
```

The system stops at the boot environment prompt.

11. Cable the port.

12. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	Reboot the system by entering the following command: system node reboot
Is a V-Series system or has FlexArray Software and is attached to storage arrays	Go to the <i>Joining the new nodes to the cluster</i> section.

13. Repeat Step 1 through Step 12 on the other node.

Decommissioning the old system

After upgrading, you can decommission the old system through the NetApp Support Site. Decommissioning the system tells NetApp that the system is no longer in operation and removes it from support databases.

Steps

- Go to the NetApp Support Site at mysupport.netapp.com and log in.
- Click the link **My Installed Systems**.
- On the **Installed Systems** page, enter the serial number of the old system in the form and then click **Go!**
A new page displays information about the controller.
- Make sure that the information about the controller is correct.

If the information about the controller is...	Then...
Correct...	<ol style="list-style-type: none"> Select Decommission this system in the Product Tool Site drop-down menu. Go to Step 5.
Not correct...	<ol style="list-style-type: none"> Click the feedback link to open the form for reporting the problem. Fill out and submit the form.

- On the **Decommission Form** page, fill out the form and click **Submit**.

How to send comments about documentation and receive update notification

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.