



Stage 3. Install and boot node3

Upgrade controllers

NetApp

February 22, 2024

Table of Contents

- Stage 3. Install and boot node3 1
 - Overview 1
 - Install and boot node3 1
 - Set the FC or UTA/UTA2 configuration on node3 12
 - Map ports from node1 to node3 20
 - Move NAS data LIFs owned by node1 from node2 to node3 and verify SAN LIFs on node3 25
 - Worksheet: Information to record before moving NAS data LIFs to node3 31
 - Relocate non-root aggregates from node2 to node3 32
 - Move NAS data LIFs owned by node2 to node3 35

Stage 3. Install and boot node3

Overview

During Stage 3, you install and boot node3, map the cluster and node-management ports from node1 to node3, and move data LIFs and SAN LIFs belonging to node1 from node2 to node3. You also relocate all aggregates from node2 to node3 and move the data LIFs and SAN LIFs owned by node2 to node3.

Steps

1. [Install and boot node3](#)
2. [Set the FC or UTA/UTA2 configuration on node3](#)
3. [Map ports from node1 to node3](#)
4. [Move NAS data LIFs owned by node1 from node2 to node3 and verify SAN LIFs on node3](#)
5. [Relocate non-root aggregates from node2 to node3](#)
6. [Move NAS data LIFs owned by node2 to node3](#)

Install and boot node3

You must install node3 in the rack, transfer node1's connections to node3, boot node3, and install ONTAP. You must also reassign any of node1's spare disks, any disks belonging to the root volume, and any non-root aggregates not relocated to node2 earlier.

About this task

You must netboot node3 if it does not have the same version of ONTAP 9 that is installed on node1. After you install node3, boot it from the ONTAP 9 image stored on the web server. You can then download the correct files to the boot media device for subsequent system boots. See [Prepare for netboot](#).

However, you do not need to netboot node3 if it has the same or a later version of ONTAP 9 that is installed on node1.



If you are upgrading a V-Series system connected to storage arrays or a system with FlexArray Virtualization software that is connected to storage arrays, you must complete [Step 1](#) through [Step 5](#), leave this section at [Step 6](#) and follow instructions in [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) as needed, entering commands in maintenance mode. You must then return to this section and resume with [Step 7](#).

However, if you are upgrading a system with storage disks, you must complete this entire section and then go to [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#), entering commands at the cluster prompt.

Steps

1. Make sure that you have rack space for node3.

If node1 and node2 were in separate chassis, you can put node3 in the same rack location as node1. However, if node1 was in the same chassis with node2, then you must put node3 into its own rack space, preferably close to the location of node1.

2. Install node3 in the rack, following the *Installation and Setup Instructions* for your node model.



If you are upgrading to a system with both nodes in the same chassis, install node4 in the chassis as well as node3. If you do not, when you boot node3, the node will behave as if it were in a dual-chassis configuration, and when you boot node4, the interconnect between the nodes will not come up.

3. Cable node3, moving the connections from node1 to node3.

The following references help you make proper cable connections. Go to [References](#) to link to them.

- *Installation and Setup Instructions* or *FlexArray Virtualization Installation Requirements and Reference* for the node3 platform
- The appropriate disk shelf procedure
- The *High Availability management* documentation

Cable the following connections:

- Console (remote management port)
- Cluster ports
- Data ports
- Cluster and node management ports
- Storage
- SAN configurations: iSCSI Ethernet and FC switch ports



You might not need to move the interconnect card or the cluster interconnect cable connection from node1 to node3 because most platform models have a unique interconnect card model. For the MetroCluster configuration, you must move the FC-VI cable connections from node1 to node3. If the new host does not have an FC-VI card, you might need to move the FC-VI card.

4. Turn on the power to node3, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.

If you are upgrading to a system with both nodes in the same chassis, node4 also reboots. However, you can disregard the node4 boot until later.



When you boot node3, you might see the following warning 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'
```


5. If you see the warning message in [Step 4](#), 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.

6. Take one of the following actions:

If your system...	Then...
Has disks and no back-end storage	Skip Step 7 through Step 12 and go to Step 13 .
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 Set the FC or UTA/UTA2 configuration on node3 and complete the subsections Configure FC ports on node3 and Check and configure UTA/UTA2 ports on node3, as appropriate to your system. b. Return to this section and complete the remaining steps, beginning with Step 7. <div>  <p>You must reconfigure FC onboard ports, CNA onboard ports, and CNA cards before you boot ONTAP on the V-Series or system with FlexArray Virtualization software.</p> </div>

7. 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.

8. 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.

9. Modify 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 values associated with each onboard FC port.

10. If your configuration uses switch-based zoning, adjust the zoning to reflect the new WWPN values.
11. Verify that the array LUNs are now visible to node3:

```
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.

12. Press Ctrl-C to display the boot menu and select maintenance mode.
13. At the Maintenance mode prompt, enter the following command:

```
halt
```

The system stops at the boot environment prompt.

14. 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 15 .
Single-chassis configuration (with controllers in the same chassis)	<div><div>a. Switch the console cable from node3 to node4.</div><div>b. Turn on the power to node4, and then interrupt the boot process by pressing Ctrl-C at the console terminal to access the boot environment prompt.</div><div>The power should already be on if both controllers are in the same chassis.</div><div><div></div><div>Leave node4 at the boot environment prompt; you will return to node4 in Install and boot node4.</div></div><div>c. If you see the warning message displayed in Step 4, follow the instructions in Step 5</div><div>d. Switch the console cable back from node4 to node3.</div><div>e. Go to Step 15.</div></div>

15. Configure node3 for ONTAP:

```
set-defaults
```

16. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps.



If you have not already done so earlier in the procedure, see the Knowledge Base article [How to tell if a drive is FIPS certified](#) to determine the type of self-encrypting drives that are in use.

- a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use...	Then...
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	<code>setenv bootarg.storageencryption.support true</code>
NetApp non-FIPS SEDs	<code>setenv bootarg.storageencryption.support false</code>



You cannot mix FIPS drives with other types of drives on the same node or HA pair.

You can mix SEDs with non-encrypting drives on the same node or HA pair.

- b. Go to the special boot menu and select option (10) Set Onboard Key Manager recovery secrets.

Enter the passphrase and the backup information that you recorded earlier procedure. See [Manage authentication keys using the Onboard Key Manager](#).

17. If the version of ONTAP installed on node3 is the same or later than the version of ONTAP 9 installed on node1, list and reassign disks to the new node3:

```
boot_ontap
```



If this new node has ever been used in any other cluster or HA pair, you must run `wipeconfig` before proceeding. Failure to do so might result in service outages or data loss. Contact technical support if the replacement controller was previously used, especially if the controllers were running ONTAP running in 7-Mode.

18. Press CTRL-C to display the boot menu.
19. Take one of the following actions:


If the system you are upgrading...	Then...
Does <i>not</i> have the correct or current ONTAP version on node3	Go to Step 20 .
Has the correct or current version of ONTAP on node3	Go to Step 25 .

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



You must use the management port and IP as the netboot connection. Do not use a data LIF IP or else a data outage might occur while the upgrade is being performed.

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Running	Configure the connection automatically by entering the following command at the boot environment prompt: <code>ifconfig e0M -auto</code>

If Dynamic Host Configuration Protocol (DHCP) is...	Then...
Not running	<p>Manually configure the connection by entering the following command at the boot environment prompt:</p> <pre>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></pre> <p><i>filer_addr</i> is the IP address of the storage system (mandatory). <i>netmask</i> is the network mask of the storage system (mandatory). <i>gateway</i> is the gateway for the storage system (mandatory). <i>dns_addr</i> is the IP address of a name server on your network (optional). <i>dns_domain</i> 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.</p> <div>  <p>Other parameters might be necessary for your interface. Enter <code>help ifconfig</code> at the firmware prompt for details.</p> </div>

21. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<pre>netboot http://<web_server_ip>/<path_to_webaccessible_directory>/netboot/kernel</pre>
All other systems	<pre>netboot http://<web_server_ip>/<path_to_webaccessible_directory>/<ontap_version>_image.tgz</pre>

The `<path_to_the_web-accessible_directory>` leads to where you downloaded the `<ontap_version>_image.tgz` in [Step 1](#) in the section *Prepare for netboot*.



Do not interrupt the boot.

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

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

Disregard the following message:

This procedure is not supported for Non-Disruptive Upgrade on an HA pair

The note applies to nondisruptive upgrades of ONTAP, and not upgrades of controllers.



Always use netboot to update the new node to the desired image. If you use another method to install the image on the new controller, the wrong image might install. This issue applies to all releases of ONTAP. The netboot procedure combined with option (7) *Install new software* wipes the boot media and places the same ONTAP version ONTAP on both image partitions.

23. If you are prompted to continue the procedure, enter `y`, and when prompted for the package, enter the following URL:

```
http://<web_server_ip>/<path_to_web-  
accessible_directory>/<ontap_version_image>.tgz
```

24. Complete the following substeps:

- a. Enter `n` to skip the backup recovery when you see the following prompt:

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

- b. 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.

25. Select **(5) Maintenance mode boot** by entering `5`, and then enter `y` when prompted to continue with the boot.
26. Before continuing, go to [Set the FC or UTA/UTA2 configuration on node3](#) to make any necessary changes to the FC or UTA/UTA2 ports on the node.

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

27. Find the system ID of node3:

```
disk show -a
```

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  HOME      DR
HOME                                NUMBER
-----
0b.02.23 nst-fas2520-2 (536880939) Pool0 KPG2RK6F nst-fas2520-
2 (536880939)
0b.02.13 nst-fas2520-2 (536880939) Pool0 KPG3DE4F nst-fas2520-
2 (536880939)
0b.01.13 nst-fas2520-2 (536880939) Pool0 PPG4KLAA nst-fas2520-
2 (536880939)
.....
0a.00.0      (536881109) Pool0 YFKSX6JG
(536881109)
.....
```



You might see the message `disk show: No disks match option -a.` after entering the command. This is not an error message so you can continue with the procedure.

28. Reassign node1's spares, any disks belonging to the root, and any non-root aggregates that were not relocated to node2 earlier in [Relocate non-root aggregates from node1 to node2](#).

Enter the appropriate form of the `disk reassign` command based on whether your system has shared disks:



If you have shared disks, hybrid aggregates, or both on your system, you must use the correct `disk reassign` command from the following table.

If disk type is...	Then run the command...
With shared disks	<code>disk reassign -s <i>node1_sysid</i> -d <i>node3_sysid</i> -p <i>node2_sysid</i></code>
Without shared disks	<code>disk reassign -s <i>node1_sysid</i> -d <i>node3_sysid</i></code>

For the `node1_sysid` value, use the information captured in [Record node1 information](#). To obtain the value for `node3_sysid`, use the `sysconfig` command.



The `-p` option is only required in maintenance mode when shared disks are present.

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)?
```

29. 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)?
```

30. 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)?
```

31. Enter `y`.

32. If you are upgrading from a system with external disks to a system that supports internal and external disks (AFF A800 systems, for example), set the node1 aggregate as root to confirm that node3 boots from the root aggregate of node1.



Warning: You must perform the following substeps in the exact order shown; failure to do so might cause an outage or even data loss.

The following procedure sets node3 to boot from the root aggregate of node1:

a. Check the RAID, plex, and checksum information for the node1 aggregate:

```
aggr status -r
```

b. Check the status of the node1 aggregate:

```
aggr status
```

c. Bring the node1 aggregate online, if necessary:

```
aggr_online root_aggr_from_node1
```

d. Prevent the node3 from booting from its original root aggregate:

```
aggr offline root_aggr_on_node3
```

- e. Set the node1 root aggregate as the new root aggregate for node3:

```
aggr options aggr_from_node1 root
```

- f. Verify that the root aggregate of node3 is offline and the root aggregate for the disks brought over from node1 is online and set to root:

```
aggr status
```



Failing to perform the previous substep might cause node3 to boot from the internal root aggregate, or it might cause the system to assume a new cluster configuration exists or prompt you to identify one.

The following shows an example of the command output:

```
-----  
      Aggr State      Status      Options  
aggr0_nst_fas8080_15 online  raid_dp, aggr  root, nosnap=on  
                        fast zeroed  
                        64-bit  
  
      aggr0 offline      raid_dp, aggr  diskroot  
                        fast zeroed  
                        64-bit  
-----
```

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

```
ha-config show
```

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 programmable ROM (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
```

```
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
```

```
ha-config modify chassis mcc
```

34. 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?
```

35. Enter `y` at the prompt to confirm that you want to destroy the local mailboxes.

36. Exit maintenance mode:

```
halt
```

The system stops at the boot environment prompt.

37. On node2, check the system date, time, and time zone:

```
date
```

38. On node3, check the date at the boot environment prompt:

```
show date
```

39. If necessary, set the date on node3:

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

40. On node3, check the time at the boot environment prompt:

```
show time
```

41. If necessary, set the time on node3:

```
set time hh:mm:ss
```

42. Verify the partner system ID is set correctly as noted in [Step 28](#) under `-p` switch:

```
printenv partner-sysid
```

43. If necessary, set the partner system ID on node3:

```
setenv partner-sysid node2_sysid
```

Save the settings:

```
saveenv
```

44. Access the boot menu at the boot environment prompt:

```
boot_ontap menu
```

45. At the boot menu, select option **(6) Update flash from backup config** by entering 6 at the prompt.

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?:
```

46. Enter `y` at the prompt.

The boot proceeds normally, and the system then asks you to confirm the system ID mismatch.



The system might reboot twice before displaying the mismatch warning.

47. 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 node might go through one round of reboot before booting normally.

48. Log in to node3.

Set the FC or UTA/UTA2 configuration on node3

If node3 has onboard FC ports, onboard unified target adapter (UTA/UTA2) ports, or a UTA/UTA2 card, you must configure the settings before completing the rest of the procedure.

About this task

You might need to complete [Configure FC ports on node3](#), or [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term "UTA2" to refer to CNA adapters and ports. However, the CLI uses the term "CNA".

- If node3 does not have onboard FC ports, onboard UTA/UTA2 ports, or a UTA/UTA2 card, and you are upgrading a system with storage disks, you can skip to the [Map ports from node1 to node3](#).
- However, if you have a V-Series system or a system with FlexArray Virtualization software with storage

arrays, and node3 does not have onboard FC ports, onboard UTA/UTA ports, or a UTA/UTA2 card, return to *Install and boot node3* and resume at [Step 22](#).

Choices:

- [Configure FC ports on node3](#)
- [Check and configure UTA/UTA2 ports on node3](#)

Configure FC ports on node3

If node3 has FC ports, either 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 from node1 that you saved in [Prepare the nodes for upgrade](#).

About this task

You can skip this section if your system does not have FC configurations. If your system has onboard UTA/UTA2 ports or a UTA/UTA2 card, you configure them in [Check and configure UTA/UTA2 ports on node3](#).



If your system has storage disks, enter the commands in this section at the cluster prompt. If you have a V-Series system or have FlexArray Virtualization Software and are connected to storage arrays, enter 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 5
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 2

2. Boot node3 and access maintenance mode:

```
boot_ontap maint
```

3. Take one of the following actions:


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

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

4. Compare the FC settings of node3 with the settings that you captured earlier from node1.
5. 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 node1	Go to Step 11 .
Different from the ones that you captured on node1	Go to Step 6 .

6. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<p>Modify the FC ports on node3 as needed by entering one of the following commands:</p> <ul style="list-style-type: none">• To program target ports: <code>system node hardware unified-connect modify -type -t target -adapter <i>port_name</i></code>• To program initiator ports: <code>system node hardware unified-connect modify -type -t initiator -adapter <i>port_name</i></code> <p>-t is the FC4 type: target or initiator.</p>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>Modify the FC ports on node3 as needed by entering the following command:</p> <pre>ucadmin modify -m fc -t initiator -f adapter_port_name</pre> <p>-t is the FC4 type, target or initiator.</p> <div> The FC ports must be programmed as initiators.</div>

7. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<p>Verify the new settings by entering the following command and examining the output:</p> <pre>system node hardware unified-connect show</pre>

If the system that you are upgrading...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Verify the new settings by entering the following command and examining the output: <code>ucadmin show</code>

8. Exit maintenance mode by entering the following command:

```
halt
```

9. After you enter the command, wait until the system stops at the boot environment prompt.
10. Take one of the following actions:

If the system you are upgrading...	Then...
Is a V-Series system or has FlexArray Virtualization software running clustered Data ONTAP 8.3	Boot node3 and access maintenance at the boot environment prompt: <code>boot_ontap maint</code>
Is not a V-Series system or does not have FlexArray Virtualization software	Boot node3 at the boot environment prompt: <code>boot_ontap</code>

11. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	<ul style="list-style-type: none"> If node3 has a UTA/UTA2 card or UTA/UTA2 onboard ports, go to Check and configure UTA/UTA2 ports on node3. If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip Check and configure UTA/UTA2 ports on node3 and go to Map ports from node1 to node3.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<ul style="list-style-type: none"> If node3 has a card or onboard ports, go to Check and configure UTA/UTA2 ports on node3. If node3 does not have a card or onboard ports, skip Check and configure UTA/UTA2 ports on node3, and return to <i>Install and boot node3</i> and resume at Step 7.

Check and configure UTA/UTA2 ports on node3

If node3 has onboard UTA/UTA2 ports or a UTA/UTA2 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 UTA/UTA2 ports.

About this task

If you want to use a Unified Target Adapter (UTA/UTA2) port for FC, you must first verify how the port is configured.



NetApp marketing materials might use the term UTA2 to refer to CNA adapters and ports. However, the CLI uses the term CNA.

You can use the `ucadmin show` command to verify the current port configuration:

```
*> ucadmin show
      Current  Current  Pending  Pending  Admin
Adapter Mode    Type    Mode    Type    Status
-----
0e      fc      target  -        initiator offline
0f      fc      target  -        initiator offline
0g      fc      target  -        initiator offline
0h      fc      target  -        initiator offline
1a      fc      target  -        -        online
1b      fc      target  -        -        online
6 entries were displayed.
```

UTA/UTA2 ports can be configured into native FC mode or UTA/UTA2 mode. FC mode supports FC initiator and FC target; UTA/UTA2 mode enables concurrent NIC and FCoE traffic sharing the same 10GbE SFP+ interface and supports FC targets.

UTA/UTA2 ports might be found on an adapter or on the controller, and have the following configurations, but you should check the configuration of the UTA/UTA2 ports on the node3 and change it, if necessary:

- UTA/UTA2 cards ordered when the controller is ordered are configured before shipment to have the personality you request.
- UTA/UTA2 cards ordered separately from the controller are shipped with the default FC target personality.
- Onboard UTA/UTA2 ports on new controllers are configured before shipment to have the personality you request.



Attention: If your system has storage disks, you must enter the commands in this section at the cluster prompt unless directed to enter maintenance mode. If you have a VSeries system or have FlexArray Virtualization Software and are connected to storage arrays, you must enter commands in this section at the maintenance mode prompt. You must be in maintenance mode to configure UTA/UTA2 ports.

Steps

1. Check how the ports are currently configured entering on of the following commands on node3:

If the system...	Then...
Has storage disks	<code>system node hardware unified-connect show</code>

If the system...	Then...
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	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

12 entries were displayed.

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin 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

```
*>
```

2. If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module.

Contact your NetApp representative to obtain the correct SFP+ module.

3. Examine the output of the `system node hardware unified-connect show` or `ucadmin show` command to determine whether the UTA/UTA2 ports have the personality you want.
4. Take one of the following actions:

If the UTA/UTA2 ports...	Then...
Do not have the personality that you want	Go to Step 5 .
Have the personality that you want	Skip Step 5 through Step 12 and go to Step 13 .

5. Take one of the following actions:

If the system...	Then...
Has storage disks and is running clustered Data ONTAP 8.3	Boot node3 and enter maintenance mode: <code>boot_ontap maint</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	Go to Step 6 . You should already be in maintenance mode.

6. Take one of the following actions:

If you are configuring...	Then...
Ports on a UTA/UTA2 card	Go to Step 7 .
Onboard UTA/UTA2 ports	Skip Step 7 and go to Step 8 .

7. If the adapter is in initiator mode, and if the UTA/UTA2 port is online, take the UTA/UTA2 port offline:

```
storage disable adapter adapter_name
```

Adapters in target mode are automatically offline in maintenance mode.

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

```
ucadmin modify -m fc|cna -t initiator|target adapter_name
```

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



You need to use the FC initiator for tape drives, FlexArray Virtualization systems, and MetroCluster configurations. You need to use the FC target for SAN clients.

9. Verify the settings:

```
ucadmin show
```

10. Verify the settings:

If the system...	Then...
Has storage disks	<p>a. Stop the system:</p> <pre>halt</pre> <p>The system stops at the boot environment prompt.</p> <p>b. Enter the following command:</p> <pre>boot_ontap</pre>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<p>Reboot to maintenance mode:</p> <pre>boot_netapp maint</pre>

11. Verify the settings:

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

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	Admin 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.

```
*> ucadmin show
```

Adapter	Current Mode	Current Type	Pending Mode	Pending Type	Admin Status
1a	fc	initiator	-	-	online
1b	fc	target	-	initiator	online
2a	fc	target	cna	-	online
2b	fc	target	cna	-	online

```
*>
```

12. Place any target ports online by entering one of the following commands, once for each port:

If the system...	Then...
Has storage disks	<code>network fcp adapter modify -node <i>node_name</i> -adapter <i>adapter_name</i> -state up</code>
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>fcp config <i>adapter_name</i> up</code>

13. Cable the port.
14. Take one of the following actions:

If the system...	Then...
Has storage disks	Go to Map ports from node1 to node3 .
Is a V-series system or has FlexArray Virtualization Software and is connected to storage arrays	Return to <i>Install and boot node3</i> and resume at Step 7 .

Map ports from node1 to node3

You must make sure that the physical ports on node1 map correctly to the physical ports on node3, which will let node3 communicate with other nodes in the cluster and with the network after the upgrade.

Before you begin

You must already have information about the ports on the new nodes from the *Hardware Universe*. (Go to [References](#) to link to the *Hardware Universe*). You use the information later in this section and in [Map ports from node2 to node4](#).

The software configuration of node3 must match the physical connectivity of node3, and IP connectivity must be restored before you continue with the upgrade.

About this task

Port settings might vary, depending on the model of the nodes.

Steps

1. Perform the following steps to verify if the setup is a two-node switchless cluster:

- a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. Verify if the setup is a two-node switchless cluster:

```
network options switchless-cluster show
```

For example:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: false/true
```

The value of this command must match the physical state of the system.

- c. Return to the administration privilege level:

```
set -privilege admin
```

2. Make the following changes:

- a. Modify ports that will be part of Cluster broadcast domain:

```
network port modify -node node_name -port port_name -mtu 9000 -ipspace
Cluster
```

This example adds Cluster port e1b on "node1":

```
network port modify -node node1 -port e1b -ipspace Cluster -mtu 9000
```

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

```
network interface migrate -vserver Vserver_name -lif lif_name -source-node
node1 -destination-node node1 -destination-port port_name
```

When all cluster LIFs are migrated and cluster communication is established, the cluster should come into quorum.

- c. Modify the home port of the Cluster LIFs:

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

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

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast
-domain Cluster -ports node1:port
```

- e. Display the health state of node1 and node3:

```
cluster show -node node1 -fields health
```

- f. Depending on the ONTAP version running on the HA pair being upgraded, take one of the following actions:

If your ONTAP version is...	Then...
9.8 to 9.11.1	Verify that the cluster LIFs are listening on port 7700: ::> network connections listening show -vserver Cluster
9.12.1 or later	Skip this step and go to Step 3 .

Port 7700 listening on cluster ports is the expected outcome as shown in the following example for a two-node cluster:

```
Cluster::> network connections listening show -vserver Cluster
Vserver Name      Interface Name:Local Port      Protocol/Service
-----
Node: NodeA
Cluster           NodeA_clus1:7700              TCP/ctlopcp
Cluster           NodeA_clus2:7700              TCP/ctlopcp
Node: NodeB
Cluster           NodeB_clus1:7700              TCP/ctlopcp
Cluster           NodeB_clus2:7700              TCP/ctlopcp
4 entries were displayed.
```

- g. For each cluster LIF that is not listening on port 7700, set the administrative status of the LIF to down and then up:

```
::> net int modify -vserver Cluster -lif cluster-lif -status-admin down; net
int modify -vserver Cluster -lif cluster-lif -status-admin up
```

Repeat substep (f) to verify that the cluster LIF is now listening on port 7700.

3. Modify the broadcast domain memberships of physical ports hosting data LIFs.
 - a. List the reachability status of all ports:

```
network port reachability show
```

- b. Repair the reachability of the physical ports, followed by VLAN ports, by running the following command on each port, one port at a time:

```
reachability repair -node node_name -port port_name
```

A warning like the following is expected. Review and enter *y* or *n* as appropriate:

WARNING: Repairing port "node_name:port" might cause it to move into a different broadcast domain, which can cause LIFs to be re-homed away from the port. Are you sure you want to continue? {y|n}:

- c. To enable ONTAP to complete the repair, wait for about a minute after running the reachability repair command on the last port.

- d. List all broadcast domains on the cluster:

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

- e. As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not correspond to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports. As required, you can delete the newly created broadcast domains if all their member ports will become member ports of the interface groups. Delete broadcast domains:

```
broadcast-domain delete -broadcast-domain broadcast_domain
```

- f. Review the interface group configuration, and as required, add or delete member ports.

Add member ports to interface group ports:

```
ifgrp add-port -node node_name -ifgrp ifgrp_port -port port_name
```

Remove member ports from interface group ports:

```
ifgrp remove-port -node node_name -ifgrp ifgrp_port -port port_name
```

- g. Delete and re-create VLAN ports as needed. Delete VLAN ports:

```
vlan delete -node node_name -vlan-name vlan_port
```

Create VLAN ports:

```
vlan create -node node_name -vlan-name vlan_port
```



Depending on the complexity of the networking configuration of the system being upgraded, you might be required to repeat Substeps (a) to (g) until all ports are placed correctly where needed.

- 4. If there are no VLANs configured on the system, go to [Step 5](#). If there are VLANs configured, restore displaced VLANs that were previously configured on ports that no longer exist or were configured on ports that were moved to another broadcast domain.

- a. Display the displaced VLANs:

```
cluster controller-replacement network displaced-vlans show
```

- b. Restore the displaced VLANs to the desired destination port:

```
displaced-vlans restore -node node_name -port port_name -destination-port
```

destination_port

- c. Verify that all displaced VLANs have been restored:

```
cluster controller-replacement network displaced-vlans show
```

- d. VLANs are automatically placed into the appropriate broadcast domains about a minute after they are created. Verify that the restored VLANs have been placed into the appropriate broadcast domains:

```
network port reachability show
```

5. Beginning with ONTAP 9.8, ONTAP will automatically modify the home ports of LIFs if the ports are moved between broadcast domains during the network port reachability repair procedure. If a LIF's home port was moved to another node, or is unassigned, that LIF will be presented as a displaced LIF. Restore the home ports of displaced LIFs whose home ports either no longer exist or were relocated to another node.

- a. Display the LIFs whose home ports might have moved to another node or no longer exist:

```
displaced-interface show
```

- b. Restore the home port of each LIF:

```
displaced-interface restore -vserver Vserver_name -lif-name LIF_name
```

- c. Verify that all LIF home ports have been restored:

```
displaced-interface show
```

When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as "ok" for all connected ports, and the status as "no-reachability" for ports with no physical connectivity. If any ports are reporting a status other than these two, repair the reachability as outlined in [Step 3](#).

6. Verify that all LIFs are administratively up on ports belonging to the correct broadcast domains.

- a. Check for any LIFs that are administratively down:

```
network interface show -vserver Vserver_name -status-admin down
```

- b. Check for any LIFs that are operationally down:

```
network interface show -vserver Vserver_name -status-oper down
```

- c. Modify any LIFs that need to be modified to have a different home port:

```
network interface modify -vserver Vserver_name -lif LIF_name -home-port  
home_port
```



For iSCSI LIFs, modification of the home port requires the LIF to be administratively down.

- d. Revert LIFs that are not home to their respective home ports:

```
network interface revert *
```

Move NAS data LIFs owned by node1 from node2 to node3 and verify SAN LIFs on node3

Before you relocate aggregates from node2 to node3, you must move the NAS data LIFs belonging to node1 that are currently on node2 from node2 to node3. You also must verify the SAN LIFs on node3.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You will verify that the LIFs are healthy and located on appropriate ports after you bring node3 online.

Steps

1. List all the NAS data LIFs not owned by node2 by entering the following command on either node and capturing the output:

```
network interface show -role data -curr-node node2 -is-home false -home-node node3
```

2. If the cluster is configured for SAN LIFs, record the SAN LIFs adapter and switch-port configuration information in this [worksheet](#) for use later in the procedure.

- a. List the SAN LIFs on node2 and examine the output:

```
network interface show -data-protocol fc*
```

The system returns output similar to the following example:

```

cluster1::> net int show -data-protocol fc*
(network interface show)

```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	-----			
svm2_cluster1	lif_svm2_cluster1_340	up/up	20:02:00:50:56:b0:39:99	cluster1-01
1b	true			
	lif_svm2_cluster1_398	up/up	20:03:00:50:56:b0:39:99	cluster1-02
1a	true			
	lif_svm2_cluster1_691	up/up	20:01:00:50:56:b0:39:99	cluster1-01
1a	true			
	lif_svm2_cluster1_925	up/up	20:04:00:50:56:b0:39:99	cluster1-02
1b	true			

4 entries were displayed.

b. List the existing configurations and examine the output:

```
fcip adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter  fc-wwpn                      switch-port
-----
cluster1-01   0a          50:0a:09:82:9c:13:38:00      ACME Switch:0
cluster1-01   0b          50:0a:09:82:9c:13:38:01      ACME Switch:1
cluster1-01   0c          50:0a:09:82:9c:13:38:02      ACME Switch:2
cluster1-01   0d          50:0a:09:82:9c:13:38:03      ACME Switch:3
cluster1-01   0e          50:0a:09:82:9c:13:38:04      ACME Switch:4
cluster1-01   0f          50:0a:09:82:9c:13:38:05      ACME Switch:5
cluster1-01   1a          50:0a:09:82:9c:13:38:06      ACME Switch:6
cluster1-01   1b          50:0a:09:82:9c:13:38:07      ACME Switch:7
cluster1-02   0a          50:0a:09:82:9c:6c:36:00      ACME Switch:0
cluster1-02   0b          50:0a:09:82:9c:6c:36:01      ACME Switch:1
cluster1-02   0c          50:0a:09:82:9c:6c:36:02      ACME Switch:2
cluster1-02   0d          50:0a:09:82:9c:6c:36:03      ACME Switch:3
cluster1-02   0e          50:0a:09:82:9c:6c:36:04      ACME Switch:4
cluster1-02   0f          50:0a:09:82:9c:6c:36:05      ACME Switch:5
cluster1-02   1a          50:0a:09:82:9c:6c:36:06      ACME Switch:6
cluster1-02   1b          50:0a:09:82:9c:6c:36:07      ACME Switch:7
16 entries were displayed
```

3. Take one of the following actions:

If node1...	Then...
Had interface groups or VLANs configured	Go to Step 4 .
Did not have interface groups or VLANs configured	Skip Step 4 and go to Step 5 .

4. Perform the following substeps to migrate any NAS data LIFs hosted on interface groups and VLANs that were originally on node1 from node2 to node3:

- Migrate any data LIFs hosted on node2 that previously belonged to node1 on an interface group to a port on node3 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

- Modify the home port and home node of the LIF in [Substep a](#) to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home-node
node3 -home-port netport|ifgrp
```

- Migrate any data LIF hosted on node2 that previously belonged to node1 on a VLAN port to a port on node3 that is capable of hosting LIFs on the same network by entering the following command, once

for each LIF:

```
network interface migrate -vserver vservice_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

- d. Modify the home port and home node of the LIFs in [Substep c](#) to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home-node
node3 -home-port netport|ifgrp
```

5. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 6 and Step 7 , skip Step 8, and complete Step 9 through Step 12 .
SAN	Disable all the SAN LIFs on the node to take them down for the upgrade: <pre>network interface modify -vserver vservice_name -lif LIF_name -home-node node_to_upgrade -home-port netport ifgrp -status-admin down</pre>

6. If you have data ports that are not the same on your platforms, add the ports to the broadcast domain:

```
network port broadcast-domain add-ports -ipSpace IPspace_name -broadcast
-domain mgmt -ports node:port
```

The following example adds port "e0a" on node "8200-1" and port "e0i" on node "8060-1" to broadcast domain "mgmt" in the IPspace "Default":

```
cluster::> network port broadcast-domain add-ports -ipSpace Default
-broadcast-domain mgmt -ports 8200-1:e0a, 8060-1:e0i
```

7. Migrate each NAS data LIF to node3 by entering the following command, once for each LIF:

```
network interface migrate -vserver vservice_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

8. Make sure that the data migration is persistent:

```
network interface modify -vserver vservice_name -lif LIF_name -home-port
netport|ifgrp -home-node node3
```

9. Confirm that the SAN LIFs are on the correct ports on node3:

- a. Enter the following command and examine its output:

```
network interface show -data-protocol iscsi|fc -home-node node3
```

The system returns output similar to the following example:

```
cluster::> net int show -data-protocol iscsi|fc -home-node node3
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	
-----	-----	-----		
vs0				
	a0a	up/down	10.63.0.53/24	node3
a0a	true			
	data1	up/up	10.63.0.50/18	node3
e0c	true			
	rads1	up/up	10.63.0.51/18	node3
e1a	true			
	rads2	up/down	10.63.0.52/24	node3
e1b	true			
vs1				
	lif1	up/up	172.17.176.120/24	node3
e0c	true			
	lif2	up/up	172.17.176.121/24	node3
e1a	true			

- b. Verify that the new and adapter and switch-port configurations are correct by comparing the output from the `fc adapter show` command with the configuration information that you recorded in the worksheet in [Step 2](#).

List the new SAN LIF configurations on node3:

```
fc adapter show -fields switch-port,fc-wwpn
```

The system returns output similar to the following example:

```
cluster1::> fcp adapter show -fields switch-port,fc-wwpn
(network fcp adapter show)
node          adapter fc-wwpn          switch-port
-----
cluster1-01 0a      50:0a:09:82:9c:13:38:00 ACME Switch:0
cluster1-01 0b      50:0a:09:82:9c:13:38:01 ACME Switch:1
cluster1-01 0c      50:0a:09:82:9c:13:38:02 ACME Switch:2
cluster1-01 0d      50:0a:09:82:9c:13:38:03 ACME Switch:3
cluster1-01 0e      50:0a:09:82:9c:13:38:04 ACME Switch:4
cluster1-01 0f      50:0a:09:82:9c:13:38:05 ACME Switch:5
cluster1-01 1a      50:0a:09:82:9c:13:38:06 ACME Switch:6
cluster1-01 1b      50:0a:09:82:9c:13:38:07 ACME Switch:7
cluster1-02 0a      50:0a:09:82:9c:6c:36:00 ACME Switch:0
cluster1-02 0b      50:0a:09:82:9c:6c:36:01 ACME Switch:1
cluster1-02 0c      50:0a:09:82:9c:6c:36:02 ACME Switch:2
cluster1-02 0d      50:0a:09:82:9c:6c:36:03 ACME Switch:3
cluster1-02 0e      50:0a:09:82:9c:6c:36:04 ACME Switch:4
cluster1-02 0f      50:0a:09:82:9c:6c:36:05 ACME Switch:5
cluster1-02 1a      50:0a:09:82:9c:6c:36:06 ACME Switch:6
cluster1-02 1b      50:0a:09:82:9c:6c:36:07 ACME Switch:7
16 entries were displayed
```



If a SAN LIF in the new configuration is not on an adapter that is still attached to the same switch-port, it might cause a system outage when you reboot the node.

- c. If node3 has any SAN LIFs or groups of SAN LIFs that are on a port that did not exist on node1 or that need to be mapped to a different port, move them to an appropriate port on node3 by completing the following substeps:

- i. Set the LIF status to "down":

```
network interface modify -vserver vservice_name -lif LIF_name -status
-admin down
```

- ii. Remove the LIF from the port set:

```
portset remove -vserver vservice_name -portset portset_name -port-name
port_name
```

- iii. Enter one of the following commands:

- Move a single LIF:

```
network interface modify -vserver vservice_name -lif LIF_name -home
-port new_home_port
```

- Move all the LIFs on a single nonexistent or incorrect port to a new port:

```
network interface modify {-home-port port_on_node1 -home-node node1
```



```
-role data} -home-port new_home_port_on_node3
```

- Add the LIFs back to the port set:

```
portset add -vserver vserver_name -portset portset_name -port-name  
port_name
```



You must move SAN LIFs to a port that has the same link speed as the original port.

10. Modify the status of all LIFs to "up" so the LIFs can accept and send traffic on the node:

```
network interface modify -home-port port_name -home-node node3 -lif data  
-status-admin up
```

11. Enter the following command on either node and examine its output to verify that LIFs have been moved to the correct ports and that the LIFs have the status of "up" by entering the following command on either node and examining the output:

```
network interface show -home-node node3 -role data
```

12. If any LIFs are down, set the administrative status of the LIFs to "up" by entering the following command, once for each LIF:

```
network interface modify -vserver vserver_name -lif LIF_name -status-admin up
```

13. Send a post-upgrade AutoSupport message to NetApp for node1:

```
system node autosupport invoke -node node3 -type all -message "node1  
successfully upgraded from platform_old to platform_new"
```

Worksheet: Information to record before moving NAS data LIFs to node3

To help verify that you have the correct configuration after moving SAN LIFs from node2 to node3, you can use the following worksheet to record the adapter and switch-port information for each LIF.

Record the LIF adapter information from the `network interface show -data-protocol fc*` command output and the switch-port information from the `fcport adapter show -fields switch-port,fc-wwpn` command output for node2.

After you complete the migration to node3, record the LIF adapter and switch-port information for the LIFs on node3 and verify that each LIF is still connected to the same switch-port.

Node2			Node3		
LIF	adapter	switch-port	LIF	adapter	switch-port

Node2			Node3		

Relocate non-root aggregates from node2 to node3

Before you can replace node2 with node4, you must send an AutoSupport message for node2 and then relocate the non-root aggregates that are owned by node2 to node3.

Steps

1. Send an AutoSupport message to NetApp for node2:

```
system node autosupport invoke -node node2 -type all -message "Upgrading node2
from platform_old to platform_new"
```

2. Verify that the AutoSupport message was sent:

```
system node autosupport show -node node2 -instance
```

The fields "Last Subject Sent:" and "Last Time Sent:" contain the message title of the last message that was sent and the time when the message was sent.

3. Relocate the non-root aggregates:
 - a. Set the privilege level to advanced:

```
set -privilege advanced
```

- b. List the aggregates that are owned by node2:

```
storage aggregate show -owner-name node2
```

- c. Start aggregate relocation:

```
storage aggregate relocation start -node node2 -destination node3 -aggregate
-list * -ndo-controller-upgrade true
```



The command locates only non-root aggregates.

- d. When prompted, enter *y*.

Relocation occurs in the background. It can take anywhere from a few seconds to a couple of minutes to relocate an aggregate. The time includes both client outage and non-outage portions. The command does not relocate any offline or restricted aggregates.

- e. Return to the admin privilege level:

```
set -privilege admin
```

4. Verify the relocation status of node2:

```
storage aggregate relocation show -node node2
```


The output will display "Done" for an aggregate after it has been relocated.



You must wait until all of the aggregates that are owned by node2 have been relocated to node3 before proceeding to the next step.

5. Take one of the following actions:

If relocation of...	Then...
All aggregates was successful	Go to Step 6 .

If relocation of...	Then...
Any aggregates failed, or was vetoed	<p>a. Display a detailed status message:</p> <pre>storage aggregate show -instance</pre> <p>You can also check the EMS logs to see the corrective action that is needed.</p> <div data-bbox="722 415 776 468">  </div> <div data-bbox="841 411 1450 474"> <p>The event log show command lists any errors that have occurred.</p> </div> <p>b. Perform the corrective action.</p> <p>c. Set the privilege level to advanced:</p> <pre>set -privilege advanced</pre> <p>d. Relocate any failed or vetoed aggregates:</p> <pre>storage aggregate relocation start -node node2 -destination node3 -aggregate-list * -ndo -controllerupgrade true</pre> <p>e. When prompted, enter y.</p> <p>f. Return to the admin privilege level:</p> <pre>set -privilege admin</pre> <p>If necessary, you can force the relocation by using one of the following methods:</p> <ul style="list-style-type: none"> • By overriding veto checks: <pre>storage aggregate relocation start -override -vetoes true -ndo-controller-upgrade</pre> <ul style="list-style-type: none"> • By overriding destination checks: <pre>storage aggregate relocation start -override -destination-checks true -ndocontroller-upgrade</pre> <p>For more information about the storage aggregate relocation commands, go to References to link to <i>Disk and aggregate management with the CLI</i> and the <i>ONTAP 9 Commands: Manual Page Reference</i>.</p>

- Verify that all of the non-root aggregates are online on node3:

```
storage aggregate show -node node3 -state offline -root false
```

If any aggregates have gone offline or have become foreign, you must bring them online, once for each

aggregate:

```
storage aggregate online -aggregate aggr_name
```

7. Verify that all of the volumes are online on node3:

```
volume show -node node3 -state offline
```

If any volumes are offline on node3, you must bring them online, once for each volume:

```
volume online -vserver Vserver-name -volume volume-name
```

8. Verify that node2 does not own any online non-root aggregates:

```
storage aggregate show -owner-name node2 -ha-policy sfo -state online
```

The command output should not display online non-root aggregates because all of the non-root online aggregates have already been relocated to node3.

Move NAS data LIFs owned by node2 to node3

After you relocate the aggregates from node2 to node3, you need to move the NAS data LIFs owned by node2 to node3.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. You must verify that the LIFs are healthy and located on the appropriate ports after you move the LIFs from node3 to node4 and bring node4 online.

Steps

1. List all the NAS data LIFs owned by node2 by entering the following command on either node and capturing the output:

```
network interface show -data-protocol nfs|cifs -home-node node2
```

The following example shows the command output for node2:

```
cluster::> network interface show -data-protocol nfs|cifs -home-node
node2
```

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

vs0					
	a0a	up/down	10.63.0.53/24	node2	a0a
true					
	data1	up/up	10.63.0.50/18	node2	e0c
true					
	rads1	up/up	10.63.0.51/18	node2	e1a
true					
	rads2	up/down	10.63.0.52/24	node2	e1b
true					
vs1					
	lif1	up/up	172.17.176.120/24	node2	e0c
true					
	lif2	up/up	172.17.176.121/24	node2	e1a
true					

- Take one of the following actions:

If node2...	Then...
Has interface groups or VLANs configured	Go to Step 3 .
Does not have interface groups or VLANs configured	Skip Step 3 and go to Step 4 .

- Take the following steps to migrate NAS data LIFs hosted on interface groups and VLANs on node2:
 - Migrate any data LIFs hosted on an interface group on node2 to a port on node3 that is capable of hosting LIFs on the same network by entering the following command, once for each LIF:

```
network interface migrate -vserver Vserver_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

- Modify the home port and home node of the LIFs in [Substep a](#) to the port and node currently hosting the LIFs by entering the following command, once for each node:

```
network interface modify -vserver Vserver_name -lif LIF_name -home-node
node3 -homeport netport|ifgrp
```

- Migrate any LIFs hosted on VLANs on node2 to a port on node3 that is capable of hosting LIFs on the same network as that of the VLANs by entering the following command, once for each LIF:

```
network interface migrate -vserver Vserver_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

- d. Modify the home port and home node of the LIFs in [Substep c](#) to the port and node currently hosting the LIFs by entering the following command, once for each LIF:

```
network interface modify -vserver Vserver_name -lif LIF_name -home-node
node3 -homeport netport|ifgrp
```

4. Take one of the following actions:

If the cluster is configured for...	Then...
NAS	Complete Step 5 through Step 8 .
SAN	Skip Step 5 through Step 8 and then complete Step 9 .
Both NAS and SAN	Complete Step 5 through Step 9 .

5. If you have data ports that are not the same on your platforms, add the ports to the broadcast domain:

```
network port broadcast-domain add-ports -ipspace IPspace_name -broadcast
-domain mgmt -ports node:port
```

The following example adds port "e0a" on node "6280-1" and port "e0i" on node "8060-1" to broadcast domain "mgmt" in the IPspace "Default":

```
cluster::> network port broadcast-domain add-ports -ipspace Default
-broadcast-domain mgmt -ports 6280-1:e0a, 8060-1:e0i
```

6. Migrate each NAS data LIF to node3 by entering the following command, once for each LIF:

```
network interface migrate -vserver Vserver_name -lif LIF_name -destination
-node node3 -destination-port netport|ifgrp
```

7. Verify that NAS LIFs have been moved to the correct ports and that the LIFs have the status of up by entering the following command on either node and examining the output:

```
network interface show -curr-node node3 -data-protocol cifs|nfs
```

8. If any LIFs are down, set the administrative status of the LIFs to "up" by entering the following command, once for each LIF:

```
network interface modify -vserver Vserver_name -lif LIF_name -status-admin up
```

9. If you have interface groups or VLANs configured, complete the following substeps:

- a. Remove the VLANs from the interface groups:

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

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

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

The system displays interface group information for the node, as shown in the following example:

```
cluster::> network port ifgrp show -node node2 -ifgrp a0a -instance
      Node: node2
Interface Group Name: a0a
Distribution Function: ip
      Create Policy: multimode_lacp
      MAC Address: MAC_address
      ort Participation: partial
      Network Ports: e2c, e2d
      Up Ports: e2c
      Down Ports: e2d
```

- c. 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
```


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.