



Stage 3. Install and boot node3

Upgrade controllers

NetApp
December 01, 2022

Table of Contents

- Stage 3. Install and boot node3 1
 - Stage 3. Install and boot node3 1
 - Install and boot node3 1
 - Set the FC or UTA/UTA2 configuration on node3 7
 - Map ports from node1 to node3 17
 - Join the quorum when a node has a different set of network ports 22
 - Verify the node3 installation 23
 - Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3 25

Stage 3. Install and boot node3

Stage 3. Install and boot node3

During Stage 3, you install and boot node3, map the cluster and node-management ports from node1 to node3, and verify the node3 installation. Non-SAN data LIFs and non-root aggregates belonging to node1 are moved from 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. [Joining the quorum when a node has a different set of network ports](#)
5. [Verify the node3 installation](#)
6. [Move non-root aggregates and NAS data LIFs owned by node1 from 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 then reassign any of node1's spare disks, any disks belonging to the root volume, and any non-root aggregates that were not relocated to node2 earlier in the process, as outlined in this section.

About this task

The relocation operation is paused at the beginning of this stage. This process is largely automated; the operation pauses to enable you to check its status. You must manually resume the operation. In addition, you must verify the SAN LIFs have successfully moved to node3.

You need to 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, by following the instructions in [Prepare for netboot](#).

Important:

- 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 need to complete [Step 1](#) through [Step 21](#), then leave this section and follow instructions in the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) sections as needed, entering commands in Maintenance mode. You must then return to this section and resume with [Step 23](#).
- If you are upgrading a system with storage disks, you need to complete this entire section and then go to the [Configure FC ports on node3](#) and [Check and configure UTA/UTA2 ports on node3](#) sections, 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 need to 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.

Cable the following connections, using the *Installation and Setup Instructions* or the *FlexArray Virtualization Installation Requirements and Reference* for the node3 platform, the appropriate disk shelf document, and *High Availability management*.

Refer to [References](#) to link to the *FlexArray Virtualization Installation Requirements and Reference* and *High Availability management*.

- 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 need to 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.



Refer to [Prepare for netboot](#).

6. 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 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>
Not running	Manually configure the connection by entering the following command at the boot environment prompt: <code>ifconfig e0M -addr=<i>filer_addr</i> -mask=<i>netmask</i> -gw=<i>gateway</i> -dns=<i>dns_addr</i> -domain=<i>dns_domain</i></code> <i>filer_addr</i> is the IP address of the storage system. <i>netmask</i> is the network mask of the storage system. <i>gateway</i> is the gateway for the storage system. <i>dns_addr</i> is the IP address of a name server on your network. <i>dns_domain</i> is the Domain Name 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. Enter <code>help ifconfig</code> at the firmware prompt for details.

7. Perform netboot on node3:

For...	Then...
FAS/AFF8000 series systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/netboot/kernel</code>
All other systems	<code>netboot http://<web_server_ip/path_to_web-accessible_directory>/<ontap_version>_image.tgz</code>

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



Do not interrupt the boot.

- 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 ONTAP releases.

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

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

- Complete the following substeps to reboot the controller module:

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

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

- Enter `y` to reboot 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 must be restored.

- Select maintenance mode 5 from the boot menu and enter `y` when you are prompted to continue with the boot.
- 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:

```
Chassis HA configuration: ha  
Controller HA configuration: ha
```



System records 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
```

```
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
```

14. Exit maintenance mode:

```
halt
```

Interrupt the autoboot by pressing `Ctrl-C` at the boot environment prompt.

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

```
date
```

16. On node3, check the date by using the following command at the boot environment prompt:

```
show date
```

17. If necessary, set the date on node3:

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

18. On node3, check the time by using the following command at the boot environment prompt:

```
show time
```

19. If necessary, set the time on node3:

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

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

```
setenv partner-sysid node2_sysid
```

- a. Save the settings:

```
saveenv
```

21. On the new node, in boot loader, the `partner-sysid` parameter must be set. For node3, `partner-sysid` must be that of node2. Verify the `partner-sysid` for node3:

```
printenv partner-sysid
```

22. Take one of the following actions:

If your system...	Description
Has disks and no back-end storage	Go to Step 23
Is a V-Series system or a system with FlexArray Virtualization software connected to storage arrays	<p>a. Go to section Setting the FC or UTA/UTA2 configuration on node3 and complete the subsections in this section.</p> <p>b. Return to this section and complete the remaining steps, beginning with Step 23.</p> <p>Important: 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>

23. 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. If required, modify the onboard ports to initiator by referring to the [Configuring FC ports on node3](#). See your storage array and zoning documentation for further instructions on zoning.

24. 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.

25. Modify the worldwide 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.

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

27. If NetApp Storage Encryption (NSE) is in use on this configuration, the `setenv bootarg.storageencryption.support` command must be set to `true`, and the `kmip.init.maxwait` variable must be set to `off` to avoid a boot loop after the node1 configuration is loaded:

```
setenv bootarg.storageencryption.support true

setenv kmip.init.maxwait off
```

28. Boot node into boot menu:

```
boot_ontap menu
```

If you do not have FC or UTA/UTA2 configuration, execute [Check and configure UTA/UTA2 ports on node3, Step 15](#) so that node3 can recognize node1's disks.

29. For a MetroCluster configuration, V-Series systems and systems with FlexArray Virtualization software connected to storage arrays, go to [Check and configure UTA/UTA2 ports on node3, Step 15](#).

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 the section [Configure FC ports on node3](#), the section [Check and configure UTA/UTA2 ports on node3](#), or both sections.



NetApp marketing materials might use the term UTA2 to refer to converged network adapter (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](#) section.
- 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 the section *Install and boot node3* and resume at [Step 23](#).

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 the section [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](#).

Important: 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.

1. Compare the FC settings on node3 with the settings that you captured earlier from node1.
2. Take one of the following actions:

If the system that you are upgrading...	Then...
Has storage disks	In maintenance mode (option 5 at boot menu), modify the FC ports on node3 as needed: <ul style="list-style-type: none"> • To program target ports: <code>ucadmin modify -m fc -t target adapter</code> • To program initiator ports: <code>ucadmin modify -m fc -t initiator adapter</code> -t is the FC4 type: target or initiator.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	In maintenance mode (option 5 at boot menu), modify the FC ports on node3 as needed: <code>ucadmin modify -m fc -t initiator -f adapter_port_name</code> -t is the FC4 type, target or initiator. Note: The FC ports must be programmed as initiators.

3. Take one of the following actions:

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

4. Exit Maintenance mode:

```
halt
```

5. Boot the system from loader prompt:

```
boot_ontap menu
```

6. After you enter the command, wait until the system stops at the boot environment prompt.
7. Select option 5 from the boot menu for maintenance mode.
8. 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 the section Check and configure UTA/UTA2 ports on node3. • If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip the section Check and configure UTA/UTA2 ports on node3, and go to the section 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 UTA/UTA2 card or UTA/UTA2 onboard ports, go to the section Check and configure UTA/UTA2 ports on node3. • If node3 does not have a UTA/UTA2 card or UTA/UTA2 onboard ports, skip the section Check and configure UTA/UTA2 ports on node3 and return to the section <i>Install and boot node3</i> at resume at Step 23.

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 allows 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 enter the commands in this section at the cluster prompt unless directed to enter Maintenance mode. If you have a V- Series system or have FlexArray Virtualization Software and are connected to storage arrays, you 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 by entering the following command on node3:

If the system...	Then...
Has storage disks	No action required.
Is a V-Series system or has FlexArray Virtualization Software and is connected to storage arrays	<code>ucadmin show</code>

The system displays output similar to the following example:

```
*> ucadmin show
      Current  Current      Pending  Pending      Admin
Adapter Mode      Type          Mode        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
*>
```

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 `ucadmin show` command and 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 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 .

6. 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.

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

- Verify the settings:

```
ucadmin show
```

- Verify the settings:

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

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

```
*> ucadmin show
      Current      Current      Pending      Pending      Admin
Adapter Mode      Type      Mode      Type      Status
-----
1a      fc      initiator -          -          online
1b      fc      target   -          initiator online
2a      fc      target   cna        -          online
2b      fc      target   cna        -          online
*>
```

- 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>

- Cable the port.
- Take one of the following actions:

If the system...	Then...
Has storage disks	Go to Map ports from node1 to node3

If the system...	Then...
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 the section at Step 23 .

13. Exit maintenance mode:

```
halt
```

14. Boot node into boot menu by running `boot_ontap menu`. If you are upgrading to an A800, go to [Step 23](#).

15. On node3, go to the boot menu and using `22/7` and select the hidden option `boot_after_controller_replacement`. At the prompt, enter `node1` to reassign the disks of node1 to node3, as per the following example.

Expand the console output example

```
LOADER-A> boot_ontap menu
...
*****
*                                     *
* Press Ctrl-C for Boot Menu. *
*                                     *
*****
.
.
Please choose one of the following:
(1) Normal Boot.
(2) Boot without /etc/rc.
(3) Change password.
(4) Clean configuration and initialize all disks.
(5) Maintenance mode boot.
(6) Update flash from backup config.
(7) Install new software first.
(8) Reboot node.
(9) Configure Advanced Drive Partitioning.
Selection (1-9)? 22/7
.
.
      (boot_after_controller_replacement)    Boot after controller upgrade
(9a)                                     Unpartition all disks and
remove their ownership information.
(9b)                                     Clean configuration and
initialize node with partitioned disks.
(9c)                                     Clean configuration and
initialize node with whole disks.
(9d)                                     Reboot the node.
(9e)                                     Return to main boot menu.

Please choose one of the following:

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



```

.
This will replace all flash-based configuration with the last backup
to
disks. Are you sure you want to continue?: yes
.
.
Controller Replacement: Provide name of the node you would like to
replace: <name of the node being replaced>
.
.
Changing sysid of node <node being replaced> disks.
Fetched sanown old_owner_sysid = 536953334 and calculated old sys id
= 536953334
Partner sysid = 4294967295, owner sysid = 536953334
.
.
.
Terminated
<node reboots>
.
.
System rebooting...
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
.
.
System rebooting...
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
Login:
...

```

16. If the system goes into a reboot loop with the message `no disks found`, this is because it has reset the ports back to the target mode and therefore is unable to see any disks. Continue with [Step 17](#) to [Step 22](#) to resolve this.
17. Press `Ctrl-C` during autoboot to stop the node at the `LOADER>` prompt.
18. At the loader prompt, enter maintenance mode:

```
boot_ontap maint
```

19. In maintenance mode, display all the previously set initiator ports that are now in target mode:

```
ucadmin show
```

Change the ports back to initiator mode:

```
ucadmin modify -m fc -t initiator -f adapter name
```

20. Verify that the ports have been changed to initiator mode:

```
ucadmin show
```

21. Exit maintenance mode:

```
halt
```



If you are upgrading from a system that supports external disks to a system that also supports external disks, go to [Step 22](#).

If you are upgrading from a system that supports external disks to a system that supports both internal and external disks, for example, an AFF A800 system, go to [Step 23](#).

22. At the loader prompt, boot up:

```
boot_ontap
```

Now, on booting, the node can detect all the disks that were previously assigned to it and can boot up as expected.

23. 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 the root aggregate to confirm that node3 boots from the root aggregate of node1. To set the root aggregate, go to the boot menu and select option 5 to enter maintenance mode.



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. Enter maintenance mode:

```
boot_ontap maint
```

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

```
aggr status -r
```

- c. Check the status of the node1 aggregate:

```
aggr status
```

- d. If necessary, bring the node1 aggregate online:

```
aggr_online root_aggr_from_node1
```

- e. Prevent the node3 from booting from its original root aggregate:

```
aggr_offline root_aggr_on_node3
```

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

```
aggr_options aggr_from_node1 root
```

- g. 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 fast zeroed 64-bit	root, nosnap=on
aggr0	offline	raid_dp, aggr fast zeroed 64-bit	diskroot

```
-----
```

Map ports from node1 to node3

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

About this task

Refer to [References](#) to link to the *Hardware Universe* to capture information about the ports on the new nodes. You will use the information later in this section.

Port settings might vary, depending on the model of the nodes. You must make the port and LIF configuration on the original node compatible with the planned use and configuration of the new node. This is because the new node replays the same configuration when it boots, which means that when you boot node3, ONTAP will try to host LIFs on the same ports that were used on node1.

Therefore, if the physical ports on node1 do not map directly to the physical ports on node3, then software configuration changes will be required to restore cluster, management, and network connectivity after the boot. In addition, if the cluster ports on node1 do not directly map to the cluster ports on node3, node3 might not automatically rejoin quorum when it is rebooted until you change the software configuration to host the cluster LIFs on the correct physical ports.

Steps

1. Record all the node1 cabling information for node1, the ports, broadcast domains, and IPspaces, in the table:

LIF	Node1 ports	Node1 IPspaces	Node1 broadcast domains	Node3 ports	Node3 IPspaces	Node3 broadcast domains
Cluster 1						
Cluster 2						
Cluster 3						
Cluster 4						
Node management						
Cluster management						
Data 1						
Data 2						
Data 3						
Data 4						
SAN						
Intercluster port						

2. Record all the cabling information for node3, the ports, broadcast domains, and IPspaces in the table.
3. Follow these steps to verify if the setup is a two-node switchless cluster:

- a. Set the privilege level to advanced:

```
cluster::> set -privilege advanced
```

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

```
cluster::> network options switchless-cluster show
```

```
cluster::*> network options switchless-cluster show

Enable Switchless Cluster: false/true
```

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

- c. Return to the administration privilege level:

```
cluster::*> set -privilege admin

cluster::>
```

4. Follow these steps to place node3 into quorum:

- a. Boot node3. See [Install and boot node3](#) to boot the node if you have not already done so.
b. Verify that the new cluster ports are in the Cluster broadcast domain:

```
network port show -node node -port port -fields broadcast-domain
```

The following example shows that port "e0a" is in the Cluster domain on node3:

```
cluster::> network port show -node _node3_ -port e0a -fields
broadcast-domain

node      port broadcast-domain
-----  ----  -----
node3     e0a  Cluster
```

- c. If the cluster ports are not in the Cluster broadcast-domain, add them with the following command:

```
broadcast-domain add-ports -ip-space Cluster -broadcast-domain Cluster -ports
node:port
```

This example adds Cluster port "e1b" on node3:

```
network port modify -node node3 -port e1b -ip-space Cluster -mtu 9000
```

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

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

This example adds Cluster port "e1b" on node4:

```
network port modify -node node4 -port e1b -ip-space Cluster -mtu 9000
```

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

```
network interface migrate -vserver Cluster -lif lif_name -source-node node3
-destination-node node3 -destination-port port_name
```

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

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

g. Remove the old ports from the Cluster broadcast domain:

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

The following command removes port "e0d" on node3:

```
network port broadcast-domain remove-ports -ipspace Cluster -broadcast  
-domain Cluster -ports node3:e0d
```

h. Verify that node3 has rejoined quorum:

```
cluster show -node node3 -fields health
```

5. Adjust the broadcast domains hosting your cluster LIFs and node-management/clustermanagement LIFs. Confirm that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF, so you might need to migrate and modify the LIFs 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. Modify a LIF's home port:

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

6. Adjust the broadcast domain membership of network ports used for intercluster LIFs 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. If there were any ports on node1 that no longer exist on node3, follow these steps to delete them:

a. Access the advanced privilege level on either node:

```
set -privilege advanced
```

b. To delete the ports:

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

c. Return to the admin level:

```
set -privilege admin
```

9. Adjust all the LIF failover groups:

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

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-  
domainwide -failover-group fg1
```

Refer to [References](#) to link to *Network Management* or the *ONTAP 9 Commands: Manual Page Reference* for more information.

10. Verify the changes on node3:

```
network port show -node node3
```

11. Each cluster LIF must be listening on port 7700. Verify that the cluster LIFs are listening on port 7700:

```
::> network connections listening show -vserver Cluster
```

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.
```

12. If necessary, 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 Step 11 to verify that the cluster LIF is now listening on port 7700.

Join the quorum when a node has a different set of network ports

The node with the new controller boots and attempts to join the cluster automatically at first; however, if the new node has a different set of network ports, you must perform the following steps to confirm that the node successfully joins the quorum.

About this task

You can use these instructions for any relevant node. Node3 is used throughout the following sample.

Steps

1. Verify that the new cluster ports are in the Cluster broadcast domain by entering the following command and checking its output:

```
network port show -node node -port port -fields broadcast-domain
```

The following example shows that port "e1a" is in the Cluster domain on node3:

```
cluster::> network port show -node node3 -port e1a -fields broadcast-
domain
node   port broadcast-domain
-----
node3  e1a  Cluster
```

2. Add the correct ports to the Cluster broadcast domain by entering the following command and checking its output:

```
network port modify -node -port -ip-space Cluster -mtu 9000
```

This example adds Cluster port "e1b" on node3:

```
network port modify -node node3 -port e1b -ip-space Cluster -mtu 9000
```

3. Migrate the cluster LIFs to the new ports, once for each LIF, using the following command:

```
network interface migrate -vserver Cluster -lif lif_name -source-node node3 -
destination-node node3 -destination-port port_name
```

4. Modify the home port of the cluster LIFs:

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

5. If the cluster ports are not in the Cluster broadcast-domain, add them by using the following command:

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


6. Remove the old ports from the Cluster broadcast domain. You can use for any relevant node. The following command removes port "e0d" on node3:

```
network port broadcast-domain remove-ports network port broadcast-domain  
remove-ports ipspace Cluster -broadcast-domain Cluster -ports node3:e0d
```

7. Verify the node has rejoined quorum:

```
cluster show -node node3 -fields health
```

8. Adjust the broadcast domains hosting your cluster LIFs and node-management/cluster management LIFs. Confirm that each broadcast domain contains the correct ports. A port cannot be moved between broadcast domains if it is hosting or is home to a LIF, so you might need to migrate and modify the LIFs 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-port
```

- d. Modify a home port of a LIF:

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

Adjust the intercluster broadcast domains and migrate the intercluster LIFs, if necessary. The data LIFs remain unchanged.

Verify the node3 installation

After you install and boot node3, you must verify that it is installed correctly. You must wait for node3 to join quorum and then resume the relocation operation.

About this task

At this point in the procedure, the operation will have paused as node3 joins quorum.

Steps

1. Verify that node3 has joined quorum:

```
cluster show -node node3 -fields health
```

2. Verify that node3 is part of the same cluster as node2 and that it is healthy:

```
cluster show
```

3. Check the status of the operation and verify that the configuration information for node3 is the same as node1:

```
system controller replace show-details
```

If the configuration is different for node3, a system disruption might occur later in the procedure.

4. Check that the replaced controller is configured correctly for the MetroCluster configuration, the MetroCluster configuration should be in healthy state and not in switch over mode. Refer to [Verify the health of the MetroCluster configuration](#).

Re-create VLANs, interface groups, and broadcast domains on node3

After you confirm that node3 is in quorum and can communicate with node2, you must re-create node1's VLANs, interface groups, and broadcast domains on node3. You must also add the node3 ports to the newly re-created broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, go to [References](#) and link to *Network Management*.

Steps

1. Re-create the VLANs on node3 using the node1 information recorded in the [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#) section:

```
network port vlan create -node node_name -vlan vlan-names
```

2. Re-create the interface groups on node3 using the node1 information recorded in the [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#) section:

```
network port ifgrp create -node node_name -ifgrp port_ifgrp_names-distr-func
```

3. Re-create the broadcast domains on node3 using the node1 information recorded in the [Relocate non-root aggregates and NAS data LIFs owned by node1 to node2](#) section:

```
network port broadcast-domain create -ipspace Default -broadcast-domain  
broadcast_domain_names -mtu mtu_size -ports  
node_name:port_name,node_name:port_name
```

4. Add the node3 ports to the newly re-created broadcast domains:

```
network port broadcast-domain add-ports -broadcast-domain  
broadcast_domain_names -ports node_name:port_name,node_name:port_name
```

Restore key-manager configuration on node3

If you are using NetApp Volume Encryption (NVE) to encrypt volumes on the system you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not restore key-manager, when you relocate the node1 aggregates from node2 to node3 by using ARL, encrypted volumes will be taken offline.

Steps

1. To synchronize encryption configuration for Onboard Key Manager, run the following command at the cluster prompt:

For this ONTAP version...	Use this command...
ONTAP 9.6 or 9.7	<code>security key-manager onboard sync</code>

For this ONTAP version...	Use this command...
ONTAP 9.5	<code>security key-manager setup -node node_name</code>

2. Enter the cluster-wide passphrase for the Onboard Key Manager.

Move non-root aggregates and NAS data LIFs owned by node1 from node2 to node3

After you verify the node3 installation and 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 that the SAN LIFs exist 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. Resume the relocation operation:

```
system controller replace resume
```

The system performs the following tasks:

- Cluster quorum check
- System ID check
- Image version check
- Target platform check
- Network reachability check

The operation pauses at this stage in the network reachability check.

2. Manually verify that the network and all VLANs, interface groups, and broadcast domains have been configured correctly.
3. Resume the relocation operation:

```
system controller replace resume
```

```
To complete the "Network Reachability" phase, ONTAP network
configuration must
be manually adjusted to match the new physical network configuration of
the
hardware. This includes assigning network ports to the correct broadcast
domains, creating any required ifgrps and VLANs, and modifying the home-
port
parameter of network interfaces to the appropriate ports. Refer to the
"Using
aggregate relocation to upgrade controller hardware on a pair of nodes
running
ONTAP 9.x" documentation, Stages 3 and 5. Have all of these steps been
manually
completed? [y/n]
```

4. Enter *y* to continue.

5. The system performs the following checks:

- Cluster health check
- Cluster LIF status check

After performing these checks, the system relocates the non-root aggregates and NAS data LIFs owned by node1 to the new controller, node3.

The system pauses once the resource relocation is complete.

6. Check the status of the aggregate relocation and NAS data LIF move operations:

```
system controller replace show-details
```

7. Verify that the non-root aggregates and NAS data LIFs have been successfully relocated to node3.

If any aggregates fail to relocate or are vetoed, you must manually relocate the aggregates, or override either the vetoes or destination checks, if necessary. See [Relocate failed or vetoed aggregates](#) for more information.

8. Verify that the SAN LIFs are on the correct ports on node3 by completing the following substeps:

a. Enter the following command and examine its output:

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

The system returns output similar to the following example:

```
cluster::> net int show -data-protocol iscsi|fcp -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. 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 Vserver_name -lif LIF_name -status  
-admin down
```

ii. Remove the LIF from the port set:

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

iii. Enter one of the following commands:

- Move a single LIF:

```
network interface modify -vserver Vserver_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 confirm that you moved SAN LIFs to a port that has the same link speed as the original port.

c. 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
```

- d. 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`:

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

- e. 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 vservice_name -lif lif_name -status-admin  
up
```

9. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check

Copyright information

Copyright © 2022 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.