
Replacing a controller module

To replace a controller module, you must remove the faulty controller module from the system and move the replaceable components, including the 10-Gb and 8-Gb SFP+ modules, from the old controller module to the new controller module. You then install the new controller module in the system chassis.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, the partner node must be able to take over the node that is being replaced (referred to in this procedure as the impaired node).

About this task

- In these procedures, *NVRAM* refers to the NVRAM8 adapter and its components.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module of the same model type; you cannot upgrade your system by just replacing the controller module.
- You cannot change any disks or disk shelves as part of this procedure.
- In this procedure, the NVRAM adapter is moved from the old controller to the new controller. Because the system ID of the controller is derived from the NVRAM adapter, the system ID does not change. Disk reassignment is not required as part of this procedure.
- The term *system* refers to FAS, AFF, V-Series, and SA (FlexCache) systems within this platform family. The procedures apply to all platforms, unless otherwise indicated, except that clustered Data ONTAP procedures do not apply to SA systems.
- It is important that you apply the commands in these steps on the correct systems:
 - The *impaired* node is the node that is being replaced.
 - The *replacement* node is the new node that is replacing the impaired node.
 - The *healthy* node is the surviving node.

Steps

1. [Pre-replacement tasks for systems in SAN configurations](#) on page 2
2. [Shutting down the target controller](#) on page 2
3. [Removing the controller module and moving the components](#) on page 6
4. [Installing the new controller module and booting the system](#) on page 12
5. [Verifying and setting the HA state of the controller module](#) on page 14
6. [Restoring the FC configuration for HA pairs](#) on page 14
7. [Verifying the system time after replacing the controller module in an HA pair](#) on page 15
8. [Installing the firmware after replacing the controller module](#) on page 16
9. [Running diagnostics tests](#) on page 17
10. [Recabling the system](#) on page 21
11. [Booting the replacement controller](#) on page 22

- 12. [Restoring licenses on the replacement node](#) on page 23
- 13. [Completing the replacement process](#) on page 24

Pre-replacement tasks for systems in SAN configurations

If you have a SAN configuration and the controller modules are in an HA pair, you must save the FC port configuration information before replacing the controller module so that you can reenter it on the new controller module. If operating in clustered Data ONTAP, you must also check that the SCSI process is in quorum with other nodes in the cluster.

Steps

1. If your system is in an HA pair, enter the following command from the impaired node to save the FC port configuration information for the impaired node (the system on which you are replacing the controller module):

If your system is operating in...	Then use this command...
7-Mode	<code>fcadmin config</code>
Clustered Data ONTAP	<code>run local fcadmin config</code>

2. Copy and save the screen display to a safe location for later reuse.
3. If your system is operating in clustered Data ONTAP, check whether the internal SCSI blade is operational and in quorum on the impaired node:

```
event log show -node impaired_node_name -messagename scsiblade.*
```

You should see messages similar to the following, indicating that the SCSI-blade process is in quorum with the other nodes in the cluster:

```
Time Node Severity Event
-----
8/13/2012 14:03:51 ssan-6240-12 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this
node established quorum with the other nodes in the cluster.
8/13/2012 14:03:51 ssan-6240-13 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this
node established quorum with the other nodes in the cluster.
8/13/2012 14:03:48 ssan-6240-14 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this
node established quorum with the other nodes in the cluster.
8/13/2012 14:03:43 ssan-6240-11 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this
node established quorum with the other nodes in the cluster.
```

If you do not see these quorum messages, check the health of the SAN processes and resolve any issues before proceeding with the replacement.

Shutting down the target controller

You shut down or take over the target controller using different procedures, depending on the storage system hardware configuration.

Choices

- [Shutting down a controller module in an HA pair](#) on page 3
- [Shutting down a stand-alone controller module](#) on page 6

Shutting down a controller module in an HA pair

To shut down a controller module, you must determine the status of the impaired node and, if necessary, take over the impaired node so that the healthy partner continues to serve data from the impaired node's storage.

About this task

As part of this procedure, you either leave the power supplies on or turn them off, depending on your configuration:

- If you have two controller modules in the same chassis, you must leave the power supplies turned on to provide power to the healthy node.
- If you have one controller module in the chassis, but it is part of an HA pair, you should turn off the power supplies in the impaired node's chassis.

Steps

1. If the system is running clustered Data ONTAP, check the status of the nodes in the cluster:

a. Change to the advance privilege level:

```
set -privilege advanced
```

b. Enter the following command at the system console of either node:

```
cluster show -epsilon *
```

The command produces output similar to the following:

Node	Health	Eligibility	Epsilon
node1	true	true	true
node2	true	true	false
node3	true	true	false
node4	true	true	false

4 entries were displayed.

Note: Epsilon must not be on a node to be replaced.

Note: In a cluster with a single HA pair, Epsilon will not be assigned to either node.

c. Take one of the following actions, depending on the result of the command:

If...	Then...
All nodes show true for both health and eligibility and Epsilon is not assigned to the impaired node.	<p>a. Exit advanced mode:</p> <pre>set -privilege admin</pre> <p>b. Proceed to Step 3.</p>

If...	Then...
All nodes show <code>true</code> for both health and eligibility and Epsilon is assigned to the impaired node	Complete the following steps to move Epsilon: <ol style="list-style-type: none"> Remove Epsilon from the node: <code>cluster modify -node node1 -epsilon false</code> Assign Epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code> Exit advanced mode: <code>set -privilege admin</code> Go to Step 3.
The impaired node shows <code>false</code> for health and is the Epsilon node.	Complete the following steps: <ol style="list-style-type: none"> Change to the advance privilege level: <code>set -privilege advanced</code> Remove Epsilon from the node: <code>cluster modify -node node1 -epsilon false</code> Assign Epsilon to a node in the cluster: <code>cluster modify -node node4 -epsilon true</code> Exit advanced mode: <code>set -privilege admin</code> Proceed to the next step.
The impaired node shows <code>false</code> for health and is not the Epsilon node.	<ol style="list-style-type: none"> Exit advanced mode: <code>set -privilege admin</code> Proceed to the next step.
Any nodes show <code>false</code> for eligibility.	<ol style="list-style-type: none"> Resolve any cluster issues as needed before continuing with this procedure. Exit advanced mode: <code>set -privilege admin</code>
Any nodes other than the impaired node show <code>false</code> for health.	<ol style="list-style-type: none"> Correct the problems that cause the health issues on the nodes before continuing with this procedure. Exit advanced mode: <code>set -privilege admin</code>

2. Check the status of the impaired node (the node you want to perform maintenance on) by entering the following command at the system console of either node:

For...	Issue the command...
7-Mode	<code>cf status</code>
Clustered Data ONTAP	<code>storage failover show</code>

3. Take one of the following actions, depending on the result of the `cf status` or `storage failover show` command:

If...	Then...
Neither node is in takeover mode	Go to the next step in this procedure.
The healthy node took over the impaired node	The impaired node is in a state where you can begin removing it from the system chassis.
The impaired node took over the healthy node	<ol style="list-style-type: none"> a. Correct the problem that caused the takeover. b. Enter the cf giveback command (7-Mode) or storage failover giveback <i>impaired_node_name</i> command (clustered Data ONTAP) from the impaired node console. c. Go back to the Step 1.

4. Take over and power down the impaired node by taking the following steps:

- a. Enter one of the following commands from the healthy node's console and wait for the takeover to complete:

For systems operating in...	Issue the command...
7-Mode	cf takeover
Clustered Data ONTAP	<ul style="list-style-type: none"> • In Data ONTAP 8.1.0 or earlier: storage failover takeover -fromnode <i>healthy_node_name</i> • In Data ONTAP 8.1.1 or later: storage failover takeover -ofnode <i>impaired_node_name</i> • In Data ONTAP 8.2 or later: storage failover takeover -ofnode <i>impaired_node_name</i>

The impaired node is taken over and then automatically reboots and displays the `waiting for giveback...` message.

- b. Wait at least two minutes after takeover of the impaired node to ensure that the takeover was completed successfully.
- c. With the impaired node showing the `waiting for giveback...` message, shut it down.

The method you use to shut down the node depends on whether remote management via a Service Processor (SP) is used and whether or not the system is in a dual-chassis or single-chassis configuration.

Is the SP configured?	Then...
Yes	Log in to the impaired node's SP and issue the following command: system power off
No, and the system is in a dual-chassis HA pair in which each controller module is in a separate chassis.	Proceed to Step 5.
No, and the system is in a single-chassis HA pair in which both controller modules are in the same chassis and share power supplies.	At the impaired node's prompt, press Ctrl-C and respond Y to halt the node.

The impaired node is now in a state where you can proceed to the next task.

5. If you are not already grounded, properly ground yourself.
6. If the system is in a dual-chassis HA pair, turn off the power supplies for the impaired node.

Shutting down a stand-alone controller module

For a stand-alone controller module, you must perform a clean shutdown to ensure that all data has been written to disk. You must also disconnect the power supplies.

Steps

1. Enter the following command from the system console of the impaired node:

If your system is configured in...	Then issue this command...
7-Mode	<code>halt</code>
Clustered Data ONTAP	<code>halt local</code>

After you issue the command and respond **Y** to halt the node, wait until the system stops at the **LOADER** prompt.

2. If you are not already grounded, properly ground yourself.
3. Turn off the power supplies, and then unplug both power cords from the power source.

Removing the controller module and moving the components

You must remove the old controller module from the chassis and move all field-replaceable components from the old controller module to the new controller module.

About this task

Attention: If the system is in an HA pair, you must wait at least two minutes after takeover of the impaired node to ensure that the takeover was successfully completed before removing the controller module.

To reduce the possibility of damage to the replaceable components, you should minimize handling by installing the components into the new controller module as soon as you remove them from the old controller module.

Note: You must also move the SFP modules from the old controller module to the new one.

Opening the system

To access components inside the controller, you must open the system.

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management tray, and then unplug the system cables and SFPs (if needed) from the controller module, and keep track of where the cables were connected.

Leave the cables in the cable management tray so that when you reinstall the cable management tray, the cables are organized.
3. Pull the cam handle downward and slide the controller module out of the system until it catches.
4. Push in the release latch on the left side of the controller module and slide the controller module out of the system and set it on a stable, antistatic surface.

Make sure that you support the bottom of the controller module with your free hand.

Moving the PCIe cards and riser card to the new controller module

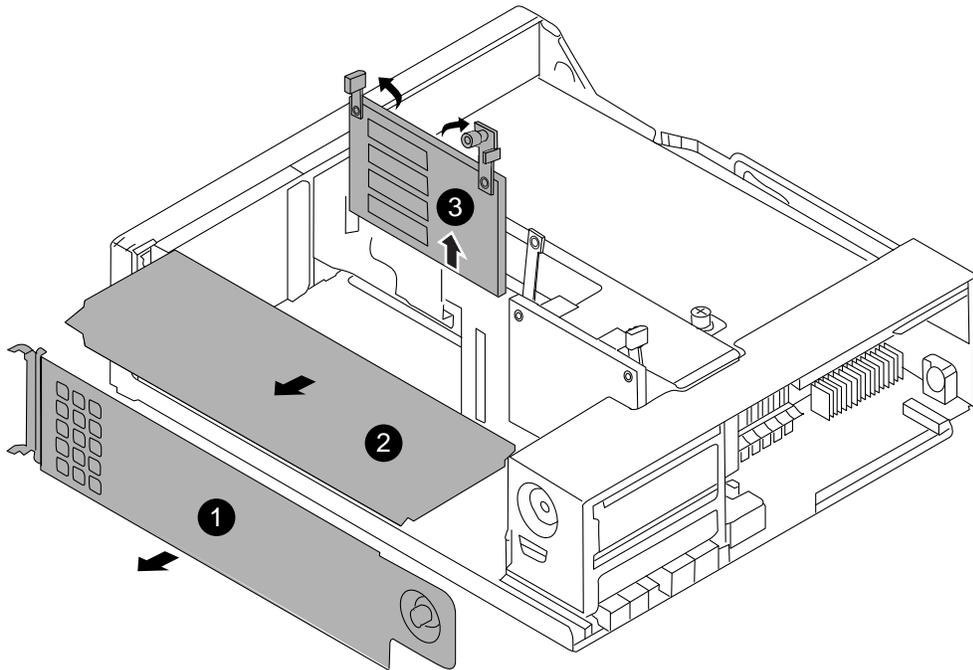
To move the PCIe cards and riser card from the old controller module to the replacement controller module, you must perform a specific sequence of steps.

Before you begin

You must have the new controller module ready so that you can move the PCIe and riser cards directly from the old controller module to the corresponding slots in the new one.

Steps

1. Loosen the thumbscrew on the controller module side panel.
2. Swing the side panel open until it comes off the controller module.



1	Side panel
2	PCIe card
3	Riser card

3. Carefully remove the PCIe card from the controller module and set it aside.
Make sure that you keep track of which slot the PCIe card was in.
4. Repeat the preceding steps for the remaining PCIe cards in the old controller module.
5. Using the FRU map inside the controller for reference, loosen the riser thumbscrew located at the top and middle of the riser, pivot the riser arms up by the tabs on the end of the arms, and then pull the riser up and out of the sockets by the tabs.

6. Install the riser into the replacement chassis by aligning the riser with the guide slots on the chassis, push the riser into the slots using the cam handles, and then rotate the handles to the closed position to seat the riser card into the connector on the controller module.
7. Tighten the riser thumbscrew.
8. Open the new controller module side panel, if necessary, slide off the PCIe card filler plate, as needed, and carefully install the PCIe card.

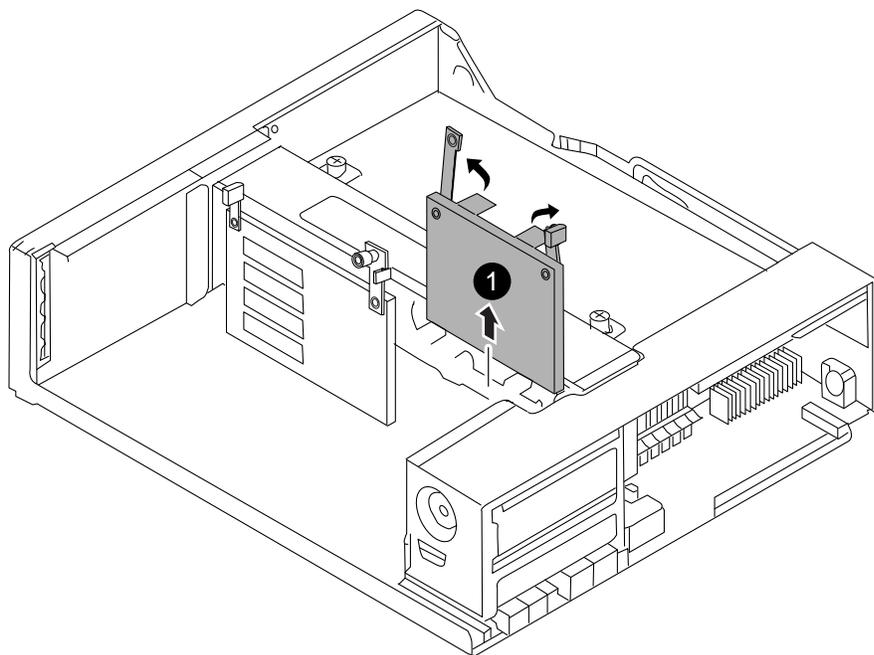
Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The card must be fully and evenly seated in the slot.
9. Repeat the preceding step as needed for additional cards.
10. Close the side panel and tighten the thumbscrew.

Moving the 10-GbE riser card

To move the 10-GbE riser card from the old controller module to the new controller module, you must perform a specific sequence of steps.

Steps

1. Loosen the thumbscrew at the top of the 10-GbE riser card and then swing the arms up.



1

10-GbE riser card

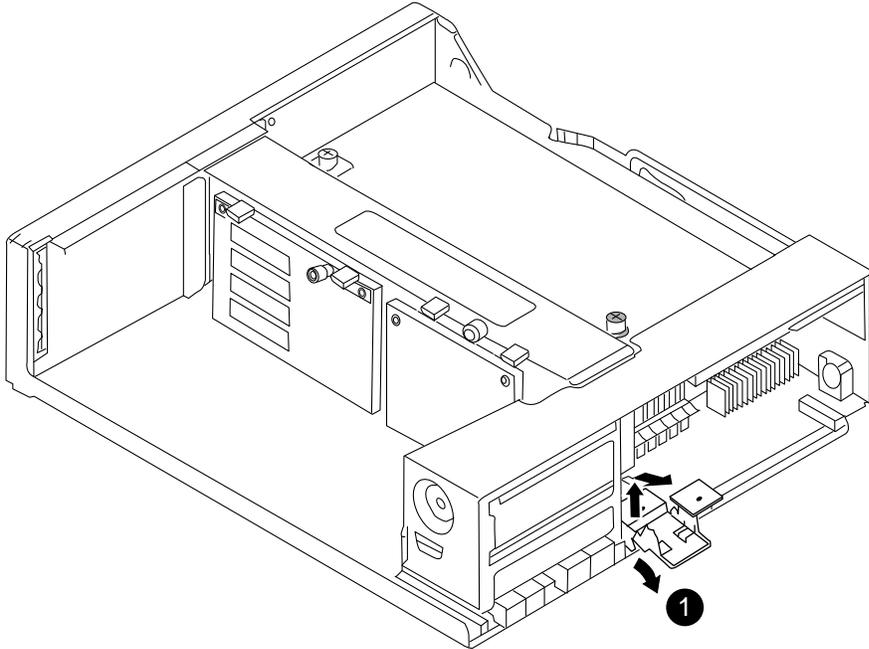
2. Pull the card straight up and lift it out of the controller module.
3. Align the 10-GbE riser card in the slot in the new controller module, push the card into the slots using the cam handles, and then rotate the handles to the closed position to seat the card into the connector on the controller module.
4. Tighten the thumbscrew on the cam handle.

Moving the boot device

To move the boot device from the old controller module to the new controller module, you must perform a specific sequence of steps.

Steps

1. Locate the boot device using the following illustration or the FRU map on the controller module:



1	Boot device
----------	-------------

2. Open the boot device cover and hold the boot device by its edges at the notches in the boot device housing, gently lift it straight up and out of the housing.
Attention: Always lift the boot device straight up out of the housing. Lifting it out at an angle can bend or break the connector pins in the boot device.
3. Open the boot device cover on the new controller module.
4. Align the boot device with the boot device socket or connector, making sure that you align the hole in the boot device with the plastic post in the boot device enclosure, and then firmly push the boot device straight down into the socket or connector.
Important: Always install the boot device by aligning the front of the boot device squarely over the pins in the socket at the front of the boot device housing. Installing the boot device at an angle or over the rear plastic pin first can bend or damage the pins in the boot device connector.
5. Check the boot device to make sure that it is seated squarely and completely in the socket or connector.
If necessary, remove the boot device and reseal it into the socket.
6. Close the boot device cover.

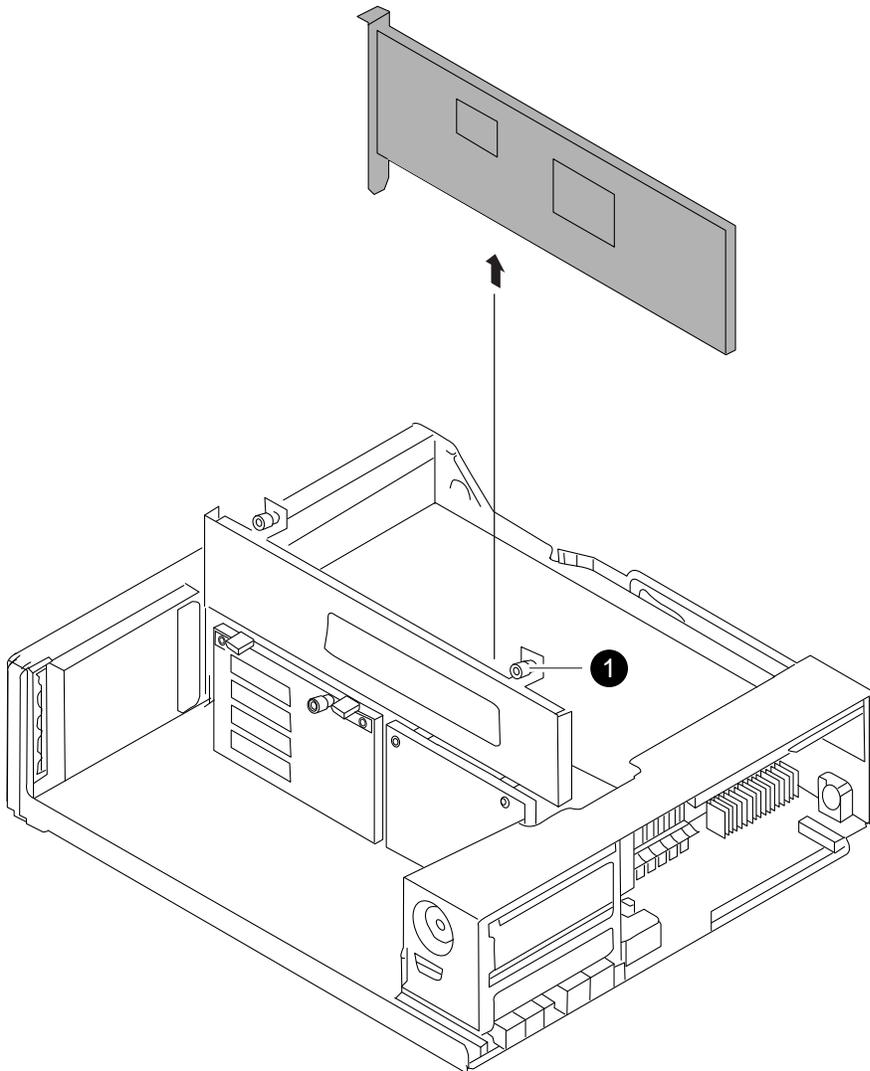
Moving the NVRAM adapter and vertical I/O card

If you want to move the NVRAM adapter and any vertical I/O card from the old controller module to the new controller module, you must perform a specific sequence of steps.

Steps

1. Loosen the two thumbscrews on the NVRAM and vertical I/O cover, and remove the cover.

The NVRAM and vertical I/O cover are located in the middle of the controller module.



1

NVRAM and vertical I/O cover thumbscrew.

There are two thumbscrews on the cover.

Note: The NVRAM8 adapter is the gray card shown lifted out of the controller module.

2. Grasp the tabs on the NVRAM adapter and lift it straight out of the controller module.
3. If present, remove the I/O card next to the NVRAM adapter.
4. Open the NVRAM and vertical I/O cover in the center of the replacement controller module, if necessary.

5. Align the NVRAM adapter in the slot (slot 2 only) and exert even pressure to seat the adapter in the socket.
6. Align the storage card in the slot and exert even pressure to seat the adapter in the socket.
Note: Slot 1 is intended for SAS or FC cards only.
7. Close the NVRAM and vertical I/O cover and tighten the cover thumbscrews.

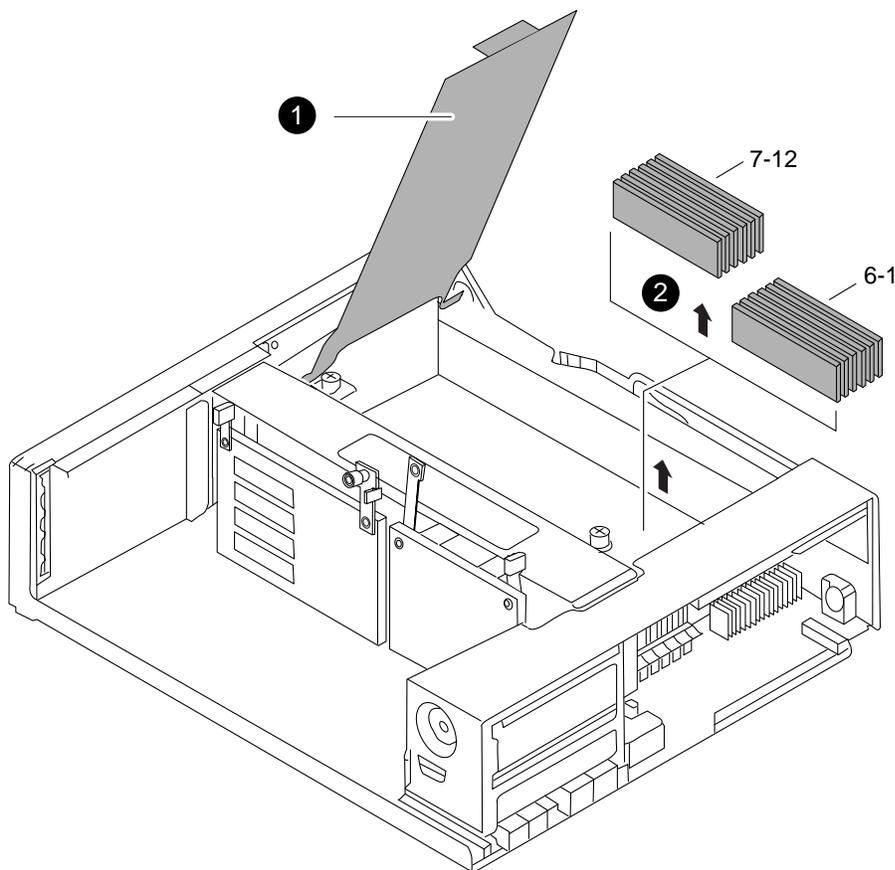
Moving the DIMMs to the new controller module

You must remove the DIMMs from the old controller module, being careful to note their locations so that you can reinstall them in the correct sockets in the new controller module.

Steps

1. Locate the DIMMs.

The number of DIMMs varies, depending on your model. This illustration shows a system fully populated with DIMMs:



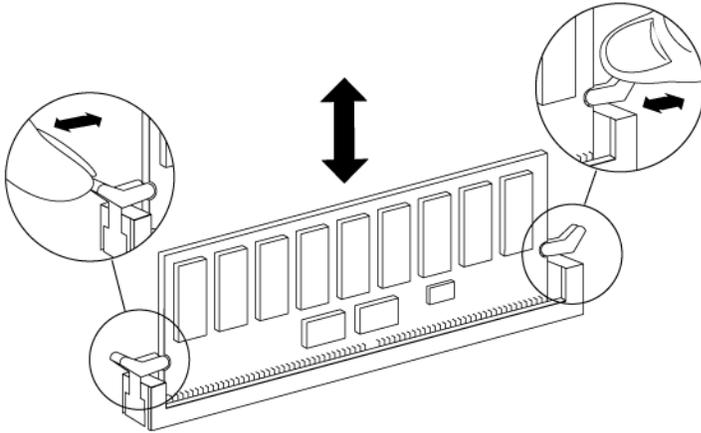
1	CPU air duct
2	<p>DIMMs</p> <p>The DIMMs are arranged in two banks, with bank 1 farthest away from the controller module cam handle containing DIMMs 1 through 6 and bank 2 containing DIMMs 7 through 12. DIMM identification numbers are silkscreened onto the motherboard.</p>

2. Note the location and orientation of the DIMM in the socket so that you can insert it in the new controller module in the proper orientation.

Note: The 6250 uses 4 GB and 8 GB DIMMs. Slots 1, 3, 5, 7, 9, and 11 use 4 GB DIMMs, while slots 2, 4, 6, 8, 10, and 12 use 8 GB DIMMs.

Check the DIMM label and the FRU map to make sure that you install the correct size DIMM in the correct slot.

3. Slowly press down on the two DIMM ejector tabs, one at a time, to eject the DIMM from its slot, and then lift it out of the slot.



Attention: Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Locate the corresponding slot for the DIMM in the new controller module, align the DIMM over the slot, and insert the DIMM into the slot.

The notch among the pins on the DIMM should align with the tab in the socket. The DIMM fits tightly in the slot but should go in easily. If not, realign the DIMM with the slot and reinsert it.

5. Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot. The edge connector on the DIMM must make complete contact with the slot.
6. Push carefully, but firmly, on the top edge of the DIMM until the latches snap into place over the notches at the ends of the DIMM.
7. Repeat these steps to move additional DIMMs, as needed.

Installing the new controller module and booting the system

After you install the components from the old controller module into the new controller module, you must install the new controller module into the system chassis and boot the operating system.

About this task

For HA pairs with two controller modules in the same chassis, the sequence in which you reinstall the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

Note: The system may update the system firmware when it boots. Do not abort this process.

Steps

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.

Note: Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the management port so you can access the system to perform the tasks in the following sections.
3. Complete the reinstall of the controller module:

If your system is in...	Then perform these steps...
An HA pair	<ol style="list-style-type: none"> a. Be prepared to interrupt the boot process. The controller module begins to boot as soon as it is fully seated in the chassis. b. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Attention: Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. c. Boot to Maintenance mode by entering halt to go to the LOADER prompt: <ul style="list-style-type: none"> • If you are running Data ONTAP 8.2.1 and earlier, enter boot_ontap, and press Ctrl-C when prompted to go to the boot menu, and then select Maintenance mode from the menu. • If you are running Data ONTAP 8.2.2 and later, enter boot_ontap maint at the LOADER prompt. d. If you have not already done so, reinstall the cable management , and then tighten the thumbscrew on the cam handle on back of the controller module. e. Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	<ol style="list-style-type: none"> a. With the cam handle in the open position, firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle to the locked position. Attention: Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors. b. Reconnect the power cables to the power supplies and to the power sources, turn on the power to start the boot process, and then press Ctrl-C to interrupt the boot process when you see the message <code>Press Ctrl-C for Boot Menu</code>. Note: If you miss the prompt and the controller module boots to Data ONTAP, enter halt and at the LOADER prompt enter boot_ontap, and press Ctrl-C when prompted, and then repeat this step. c. From the boot menu, select the option for Maintenance mode. d. If you have not already done so, reinstall the cable management , and then tighten the thumbscrew on the cam handle on back of the controller module. e. Bind the cables to the cable management device with the hook and loop strap.

Important: During the boot process, you might see the following prompts:

- A prompt warning of a system ID mismatch and asking to override the system ID.
- A prompt warning that when entering Maintenance mode in a HA configuration you must ensure that the healthy node remains down.

You can safely respond **Y** to these prompts.

Verifying and setting the HA state of the controller module

You must verify the **HA** state of the controller module and, if necessary, update the state to match your system configuration (HA pair or stand-alone).

Steps

1. In Maintenance mode, enter the following command from the new controller module to display the **HA** state of the new controller module and chassis:

```
ha-config show
```

The **HA** state should be the same for all components.

If your system is...	The HA state for all components should be...
In an HA pair	ha
Stand-alone	non-ha

2. If the displayed system state of the controller does not match your system configuration, set the **HA** state for the controller module:

```
ha-config modify controller ha-state
```

If your system is...	Issue the following command...
In an HA pair	ha-config modify controller ha
Stand-alone	ha-config modify controller non-ha

3. If the displayed system state of the chassis does not match your system configuration, set the **HA** state for the chassis:

```
ha-config modify chassis ha-state
```

If your system is...	Issue the following command...
In an HA pair	ha-config modify chassis ha
Stand-alone	ha-config modify chassis non-ha

Restoring the FC configuration for HA pairs

Because the onboard Fibre Channel (FC) ports are not preconfigured, you must restore any FC port configurations in your HA pair before you bring the node back into service; otherwise, you might experience a disruption in service. Systems without FC configurations can skip this procedure.

Before you begin

You must have the values of the FC port settings that you saved earlier.

Steps

1. From the healthy node, verify the values of the FC configuration on the replacement node:

If your system is operating in...	Then issue this command...
7-Mode or Data ONTAP versions prior to 8.0	partner fcadmin config

If your system is operating in...	Then issue this command...
Clustered Data ONTAP	<code>system node run -node healthy_node_name partner fcadmin config</code>

2. Compare the default FC variable settings with the list you saved earlier.

If the FC variables are...	Then...
The same as you recorded earlier	Proceed to the next step in this procedure.
Different than you recorded earlier	<ol style="list-style-type: none"> a. If you have not already done so, reboot the replacement node to Maintenance mode by pressing Ctrl-C when you see the message <code>Press Ctrl-C for Boot Menu</code>. b. Answer y when prompted by the system. c. Select the Maintenance mode option from the displayed menu. d. Enter one of the following commands, depending on what you need to do: <ul style="list-style-type: none"> • To program target ports: <code>fcadmin config -t target adapter_name</code> • To program initiator ports: <code>fcadmin config -t initiator adapter_name</code> • To unconfigure ports: <code>fcadmin config -t unconfigure adapter_name</code> e. Verify the values of the variables by entering the following command: <code>fcadmin config</code> f. Exit Maintenance mode by entering the following command: <code>halt</code> After you issue the command, wait until the system stops at the LOADER prompt. g. If you are operating in clustered Data ONTAP 8.2.2 and earlier, enter the following command at the LOADER prompt: <code>printenv</code> The following line should be present in the displayed list: <code>bootarg.init.boot_clustered true</code> If it is not, you should issue the following command to set the variable: <code>setenv bootarg.init.boot_clustered true</code>

Verifying the system time after replacing the controller module in an HA pair

If your system is in an HA pair, you must set the time on the replacement node to that of the healthy node to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement node* is the new node that replaced the impaired node as part of this procedure.
- The *healthy node* is the HA partner of the replacement node.

When setting the date and time at the LOADER prompt, be sure that all times are set to GMT.

Steps

1. If you have not already done so, halt the replacement node to display the LOADER prompt.
2. On the healthy node, check the system time:
`date`
3. At the LOADER prompt, check the date and time on the replacement node:
`show date`
The date and time are given in GMT.
4. If necessary, set the date in GMT on the replacement node:
`set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node:
`set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the replacement node:
`show date`
The date and time are given in GMT.

Installing the firmware after replacing the controller module

After replacing the controller module, you must install the latest firmware, and check and update the Service Processor (SP) firmware if needed, on the new controller module. If the system is in an HA pair, the healthy node should also be updated so that each controller module is running the same firmware version.

About this task

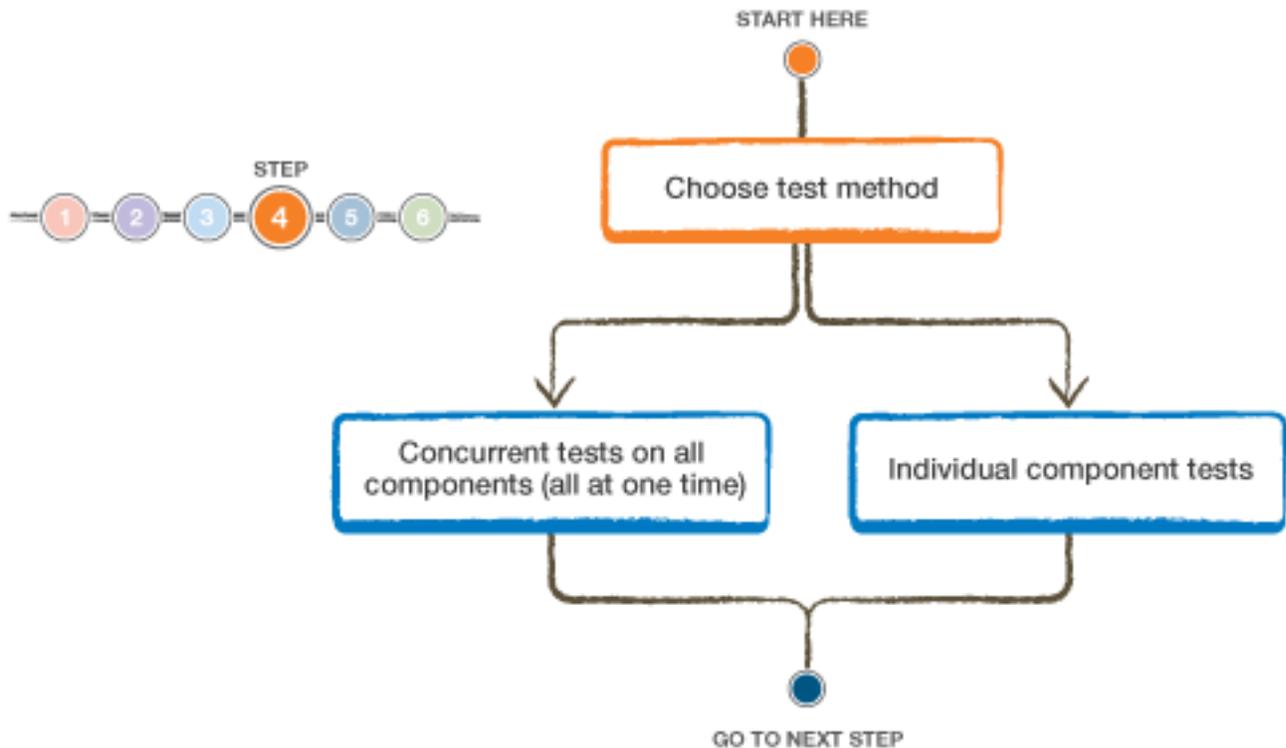
The SP firmware and BIOS automatically update, if necessary, to the baseline image included with the Data ONTAP version. Other system firmware from the old controller module still resides on the boot device and typically does not need updating.

Steps

1. Check the configuration of the SP from the LOADER prompt:
`sp status`
Note the version of the SP firmware and cross-check it and update it, if needed, when you log in to [NetApp Support](#) in the following steps.
2. Log in to the SP from an administration host to install the firmware:
`ssh username@SP_IP_address`
3. Download and install the most current version of firmware for your system, following the provided instructions.
[NetApp Downloads: System Firmware and Diagnostics](#)
Note: You can also take this opportunity to download and install the SP firmware and BIOS on the healthy node, if needed.

Running diagnostics tests

You should run focused diagnostic tests for specific components and subsystems whenever you replace a component in your system.



Before you begin

- Your system must be at the LOADER prompt to start System Level Diagnostics.
- For Data ONTAP 8.2 and later, you do not need loopback plugs to run tests on storage interfaces.

About this task

All commands in the diagnostic procedures are issued from the node where the component is being replaced.

Steps

1. If the node to be serviced is not at the LOADER prompt, bring it to the LOADER prompt.
2. On the node with the replaced component, enter the following command at the LOADER prompt:

```
boot_diags
```

Note: You must enter this command from the LOADER prompt for system-level diagnostics to function properly. The `boot_diags` command starts special drivers designed specifically for system-level diagnostics.

Important: During the `boot_diags` process, you might see the following prompt:

- A prompt warning that when entering Maintenance mode in an HA configuration you must ensure that the partner remains down.

You can safely respond **y** to the prompt.

The Maintenance mode prompt (*>) appears.

3. Clear the status logs by entering the following command:

```
sldiag device clearstatus
```

4. Display and note the available devices on the controller module by entering the following command:

```
sldiag device show -dev mb
```

The controller module devices and ports displayed can be any one or more of the following:

- **bootmedia** is the system booting device.
- **env** is motherboard environmentals.
- **mem** is system memory.
- **nic** is a network interface card.
- **nvr** is nonvolatile RAM.
- **sas** is a Serial Attached SCSI device not connected to a disk shelf.

5. How you proceed depends on how you want to run diagnostics on your system.

If you want to run...	Then go here...
Diagnostic tests on individual components	Running diagnostics tests individually on page 19
Diagnostic tests on multiple components at the same time	Running diagnostics tests concurrently on page 18

Choices

- [Running diagnostics tests concurrently](#) on page 18
- [Running diagnostics tests individually](#) on page 19

Related information

[Clustered Data ONTAP 8.3 System Administration Guide](#)

Running diagnostics tests concurrently

After replacing the controller module, you can choose to run diagnostics tests concurrently if you want a single organized log of all the test results for all the devices.

About this task

The time required to complete this procedure can vary based on the choices you make. If you run more than the default tests, the diagnostic test process will take longer to complete.

Steps

1. Review the enabled and disabled devices in the output from Step 4 in the preceding procedure and determine which ones you want to run concurrently.

2. List the individual tests for the device:

```
sldiag device show -dev dev_name
```

3. Examine the output and enable the tests that you want to run for the device by entering the following command:

```
sldiag device modify -dev dev_name -index test_index_number -selection enable
```

test_index_number can be an individual number, a series of numbers separated by commas, or a range of numbers.

4. Examine the output and disable the tests that you do not want to run for the device by entering the following command, which enables only the tests that you want to run:

```
sldiag device modify -dev dev_name -index test_index_number -selection disable
```

5. Verify that the tests were modified:

```
sldiag device show
```

6. Repeat steps 1 through 5 for each device.

7. Run diagnostics on all the devices:

```
sldiag device run
```

Attention: Do not add to or modify your entries after you start running diagnostics.

The test is complete when the following message is displayed:

```
*> <SLDIAG: _ALL_TESTS_COMPLETED>
```

8. After the test is complete, verify that there are no hardware problems on your storage system:

```
sldiag device status -long -state failed
```

9. Correct any issues found and repeat this procedure.

Result

You have completed the system-level diagnostics for the selected tests.

Running diagnostics tests individually

After replacing the controller module, you can run diagnostics tests individually if you want a separate log of all the test results for each device.

Steps

1. Clear the status logs by entering the following command:

```
sldiag device clearstatus
```

2. Display the available tests for the selected devices:

Device type	Command
boot media	sldiag device show -dev bootmedia
env	sldiag device show -dev env
mem	sldiag device show -dev mem
nic	sldiag device show -dev nic

Device type	Command
<code>nvr</code>	<code>sldiag device show -dev nvr</code>
<code>sas</code>	<code>sldiag device show -dev sas</code>

The available device tests are displayed.

- Examine the output and, if applicable, enable the tests that you want to run for the device:

```
sldiag device modify -dev dev_name -index test_index_number -selection enable
```

`test_index_number` can be an individual number, a series of numbers separated by commas, or a range of numbers.

- Examine the output and, if applicable, disable the tests that you do not want to run for the device by selecting only the tests that you want to run:

```
sldiag device modify -dev dev_name -index test_index_number -selection only
```

- Run the selected tests:

Device type	Command
<code>boot media</code>	<code>sldiag device run -dev bootmedia</code>
<code>env</code>	<code>sldiag device run -dev env</code>
<code>mem</code>	<code>sldiag device run -dev mem</code>
<code>nic</code>	<code>sldiag device run -dev nic</code>
<code>nvr</code>	<code>sldiag device run -dev nvr</code>
<code>sas</code>	<code>sldiag device run -dev sas</code>

After the test is complete, the following message is displayed:

```
<SLDIAG: _ALL_TESTS_COMPLETED>
```

- Verify that no tests failed:

Device type	Command
<code>boot media</code>	<code>sldiag device status -dev bootmedia -long -state failed</code>
<code>env</code>	<code>sldiag device status -dev env -long -state failed</code>
<code>mem</code>	<code>sldiag device status -dev mem -long -state failed</code>
<code>nic</code>	<code>sldiag device status -dev nic -long -state failed</code>
<code>nvr</code>	<code>sldiag device status -dev nvr</code>
<code>sas</code>	<code>sldiag device status -dev sas -long -state failed</code>

Any tests that failed are displayed.

- Proceed based on the result of the preceding step:

If the system-level diagnostics tests...	Then...
--	---------

Were completed without any failures

- a. Clear the status logs by entering the following command:
`sldiag device clearstatus`
 - b. Verify that the log is cleared by entering the following command:
`sldiag device status`
The following default response is displayed:
`SLDIAG: No log messages are present.`

You have completed system-level diagnostics.
-

Resulted in some test failures

- Determine the cause of the problem:
- a. Exit Maintenance mode by entering the following command:
`halt`
After you issue the command, wait until the system stops at the LOADER prompt.
 - b. Turn off or leave on the power supplies, depending on how many controller modules are in the chassis:
 - If you have two controller modules in the chassis, leave the power supplies turned on to provide power to the other controller module.
 - If you have one controller module in the chassis, turn off the power supplies and unplug them from the power sources.
 - c. Check the controller module you are servicing and verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system.
 - d. Boot the controller module you are servicing, interrupting the boot by pressing **Ctrl-C** when prompted. This takes you to the Boot Menu:
 - If you have two controller modules in the chassis, fully seat the controller module you are servicing in the chassis.
The controller module boots up when fully seated.
 - If you have one controller module in the chassis, connect the power supplies and turn them on.
 - e. Select Boot to maintenance mode from the menu.
 - f. Exit Maintenance mode by entering the following command:
`halt`
After you issue the command, wait until the system stops at the LOADER prompt.
 - g. Enter `boot_diags` at the prompt and rerun the system-level diagnostic test.
-

8. Continue to the next device that you want to test or exit system-level diagnostics and continue with the procedure.

Recabling the system

After running diagnostics, you must recable the controller module's storage and network connections. If you have a dual-chassis HA pair, you must recable the HA interconnect.

Steps

1. Recable the system, as needed.

If you removed the media converters (SFPs), remember to reinstall them if you are using fiber optic cables.

2. Check your cabling using Config Advisor.
 - a. Download and install Config Advisor from the NetApp Support Site at mysupport.netapp.com.
 - b. Enter the information for the target system and click **Collect Data**.
 - c. Click the **Cabling** tab and examine the output.

Make sure that all disk shelves are displayed and all disks appear in the output. Correct any cabling issues you might find.

- d. Check other cabling by clicking the appropriate tab and examining the output from Config Advisor.

Booting the replacement controller

After recabling the replacement controller, you can boot it up and resume normal operation.

Steps

1. Exit Maintenance mode from the *replacement* node to display the LOADER prompt:

```
halt
```

2. After the replacement node displays the LOADER prompt, boot the operating system:

```
boot_ontap
```

For a system in an HA pair, this command puts the node in *Waiting for Giveback* state.

For a system in a stand-alone configuration, the replacement node fully boots up.

3. If the system is in an HA pair:

- a. Return storage to the replacement node:

If the controller module is operating in...	Then issue this command...
7-Mode	<code>cf giveback</code>
Clustered Data ONTAP	<ul style="list-style-type: none">• In Data ONTAP 8.1 or earlier: <code>storage failover giveback -fromnode healthy_node_name</code>• In Data ONTAP 8.1.1 or later: <code>storage failover giveback -ofnode replacement_node_name</code>

- b. After the giveback operation is complete, check that the HA pair is healthy and takeover is possible:

If the controller module is operating in...	Then issue this command...
7-Mode	<code>cf show</code>
Clustered Data ONTAP	<code>storage failover show</code>

- c. Verify that the disks or V-Series LUNs are assigned correctly:

If the controller module is operating in...	Then issue this command...
7-Mode	<code>disk show -v</code>
Clustered Data ONTAP	<code>storage disk show -ownership</code>

Restoring licenses on the replacement node

You might have to install licenses on the new replacement node. The procedure varies depending on whether you are running Data ONTAP 7-Mode or clustered Data ONTAP.

Choices

- [Installing licenses for the replacement node operating in 7-Mode](#) on page 23
- [Installing licenses for the replacement node in clustered Data ONTAP](#) on page 24

Installing licenses for the replacement node operating in 7-Mode

You must reinstall new license keys for each feature package that was on the impaired node. The same license packages should be installed on both controller modules in an HA pair. Each controller module needs its own license keys.

About this task

Some features require that you enable certain options instead of, or in addition to, installing a license key. For detailed information about licensing, see the knowledgebase article 3013749: Data ONTAP 8.2 Licensing Overview and References on the NetApp Support Site and the *Data ONTAP System Administration Guide for 7-Mode*.

The licenses keys must be in the 28-character format used by Data ONTAP 8.2.

You have a 90-day grace period in which to install the license keys; after the grace period, all old licenses are invalidated. Once a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

You can use the `license show` command to check the time available before the grace period expires.

Steps

1. If you need new license keys in the Data ONTAP 8.2 format, obtain replacement license keys on the NetApp Support Site in the **My Support** section under Software licenses.

Note: The new license keys that you require are auto-generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, contact technical support.

2. Wait until the Data ONTAP command line interface has been up for at least five minutes to ensure the license database is running.

3. Issue the following command to install the license keys:

```
license add license_key license_key license_key...
```

You can add one license or multiple licenses at the same time, each license key separated by a comma or a space.

You might get a message indicating that the license database is unavailable if the Data ONTAP command line interface was not up for a sufficient amount of time.

4. Verify that the licenses have been installed:

```
license show
```

Installing licenses for the replacement node in clustered Data ONTAP

You must install new licenses for the replacement node if the impaired node was using Data ONTAP features that require a standard (node-locked) license. For features with standard licenses, each node in the cluster should have its own key for the feature.

About this task

Until you install license keys, features requiring standard licenses will continue to be available to the replacement node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the replacement node as soon as possible.

The licenses keys must be in the 28-character format used by Data ONTAP 8.2 and later versions of Data ONTAP.

You have a 90-day grace period in which to install the license keys; after the grace period, all old licenses are invalidated. Once a valid license key is installed, you have 24 hours to install all of the keys before the grace period ends.

Steps

1. If you need new license keys in the Data ONTAP 8.2 format, obtain replacement license keys on the NetApp Support Site in the My Support section under Software licenses.

Note: The new license keys that you require are auto-generated and sent to the email address on file. If you fail to receive the email with the license keys within 30 days, contact technical support.

2. Issue the following command to install each license key:

```
system license add -license-code license-key, license-key...
```

3. If you want to remove the old licenses, complete the following substeps:

- a. Check for unused licenses by entering the following command:

```
license clean-up -unused -simulate
```

- b. If the list looks correct, remove the unused licenses by entering the following command:

```
license clean-up -unused
```

Completing the replacement process

After you replace the part, you can return the failed part to NetApp, as described in the RMA instructions shipped with the kit. Contact technical support at [NetApp Support](#), 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Disposing of batteries

Dispose of batteries according to local regulations regarding battery recycling or disposal. If you cannot properly dispose of the battery, return it to NetApp, as described in the RMA instructions shipped with the kit.

Related information

[Safety Information and Regulatory Notices at support.netapp.com](#)

Copyright information

Copyright © 1994–2015 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

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

How to send comments about documentation and receive update notifications

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

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

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

You can also contact us in the following ways:

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