Replacing the controller module

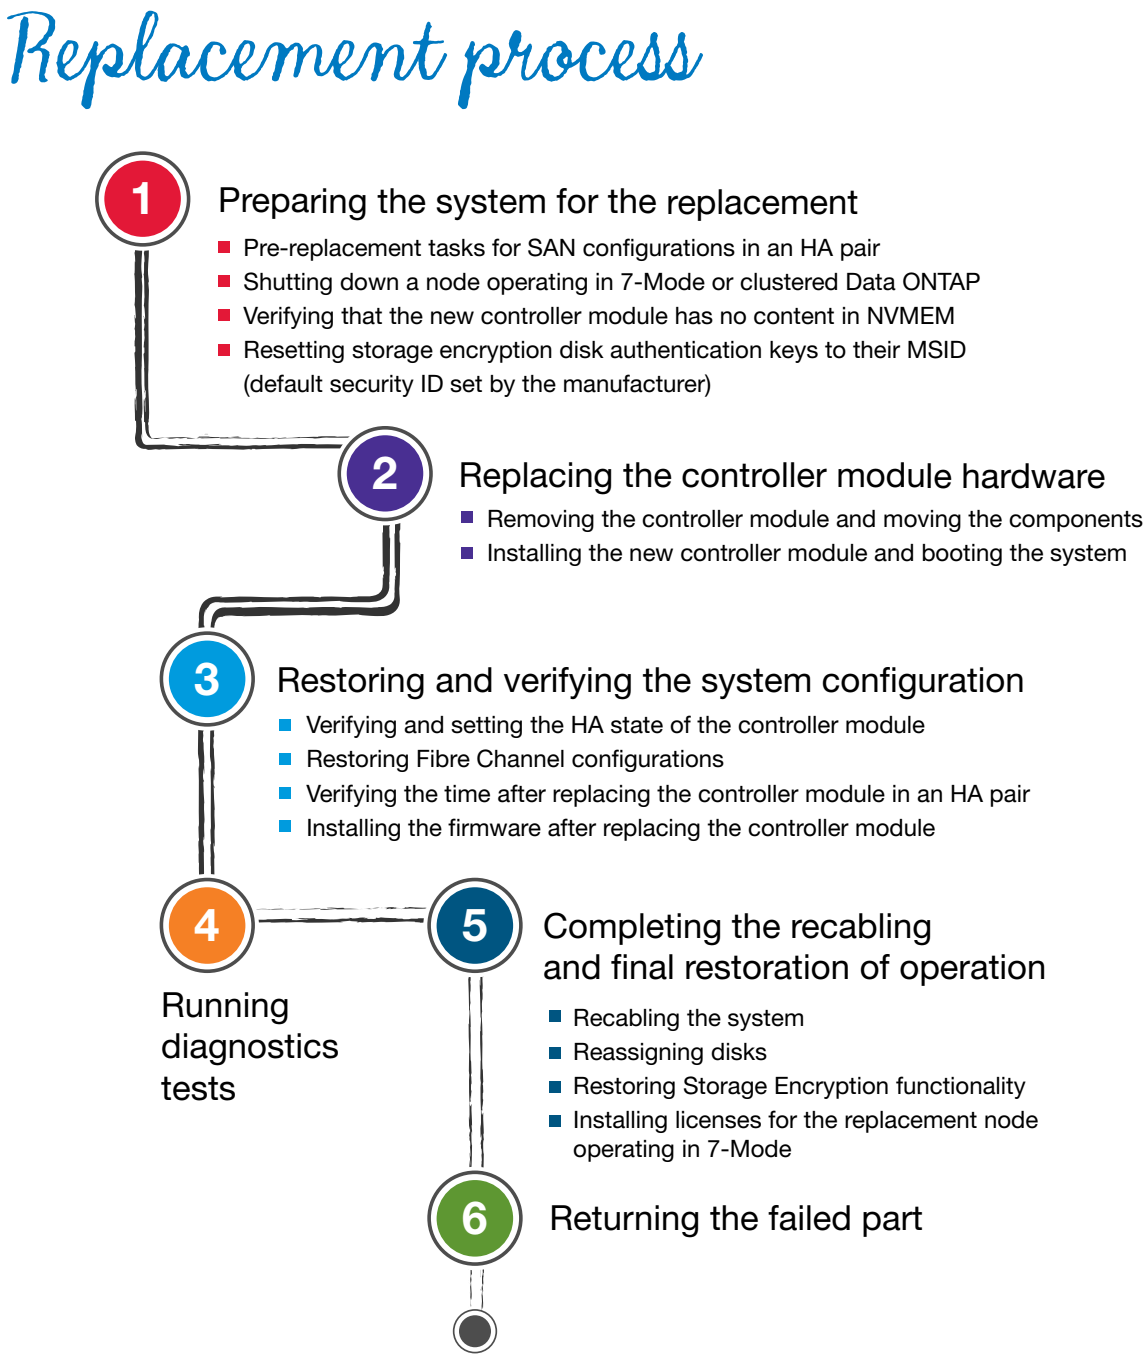

Replacing the controller module

You must review the prerequisites for the replacement procedure and then select the correct one for your Data ONTAP operating system.

Before you begin

- All disk shelves must be working properly.
- If your system is in an HA pair, then the healthy 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

- This procedure includes steps for automatically or manually reassigning disks to the replacement node, depending on your system's configuration.
  You should perform the disk reassignment as directed in the procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must replace a controller module with 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 boot device is moved from the impaired node to the replacement node so that the replacement node boots up in the same version of ONTAP as the old controller module.
- Any PCIe cards moved from the old controller module to the new controller module or added from existing customer site inventory must be supported by the replacement controller module.
  NetApp Hardware Universe
- 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.
- You must always capture the node's console output to a text file.
  This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.
- The CNA ports are referred to as UTA2 ports.

Choices

- *Replacing a controller in 7-Mode* on page 2
- *Replacing a controller in clustered Data ONTAP* on page 31

Replacing a controller module in 7-Mode environments

You must follow a specific series of steps to replace the for your mode and version of ONTAP.

Steps

1. Preparing the system for the replacement on page 3
2. Replacing the controller module hardware on page 8
3. Restoring and verifying the system configuration after hardware replacement on page 16
4. Running diagnostics tests after replacing a controller module on page 19
5. Completing the recabling and final restoration of operations on page 24
6. Completing the replacement process on page 30

Preparing the system for the replacement

You must gather information and shut down the impaired node by taking it over if it is in an HA pair.

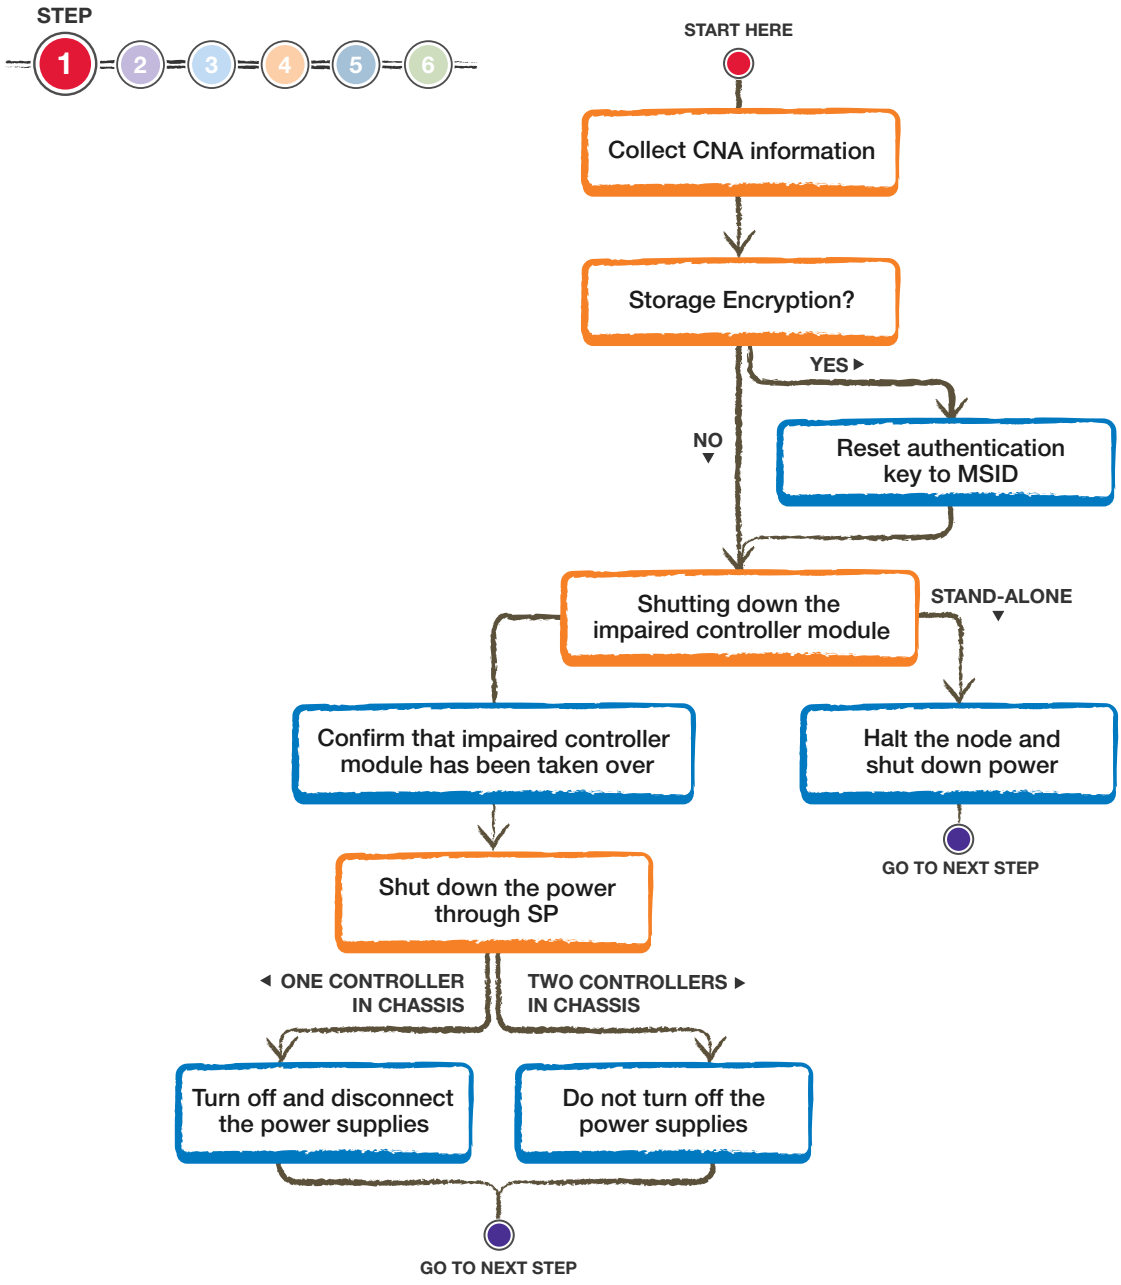

Determining your controller CNA port configuration

If you have a SAN configuration, you must save the FC port configuration information of the impaired node so that you can reenter it on the replacement node.

About this task

Your system configuration determines your access to port configuration information.

Step

1. Take one of the following actions, depending on your configuration:

If the system is in...	Then...
A stand-alone configuration and is not running	You have to rely on any configuration backups or information gathered previously from the AutoSupport tool.
An HA pair and the impaired node has not been taken over by the healthy node and is running	**a.** To save the port configuration information for the impaired node: **`ucadmin show`** **b.** Copy and save the screen display to a safe location for later reuse. **Note:** If the impaired node is taken over by its partner, you can boot it to Maintenance mode and run the `ucadmin show` command in Maintenance mode. To boot the impaired node to Maintenance mode, restart the impaired node, press Ctrl-C to interrupt the boot process when you see the message `Press Ctrl-C` for the Boot Menu. From the Boot Menu, enter the option for Maintenance mode. **c.** Enter the `Cluster-Mode` command to save the port configuration information for the impaired node: **`unified-connect modify`**

Pre-replacement tasks for Storage Encryption configurations

If the storage system whose controller you are replacing is configured to use Storage Encryption, you must first reset the authentication keys of the disks to their MSID (the default security ID set by the manufacturer). This is a temporary necessity during the controller replacement process to avoid any chance of losing access to the data.

About this task

After resetting the authentication keys to the MSID, the data on the disks is no longer encrypted with secret authentication keys. You must verify the physical safety of the disks during the replacement or upgrade process.

Steps

1. Display the key ID for each self-encrypting disk on the original system:

   `disk encrypt show`

Example

```
disk encrypt show
Disk       Key ID                                                           Locked?
0c.00.1    0x0                                                              No
0c.00.0    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
```

```
0c.00.3    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.4    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.2    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
0c.00.5    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
```

The first disk in the example is associated with an MSID; the others are associated with a non-MSID.

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

   - Rekey the disks individually, once for each disk:

     disk encrypt rekey 0x0 *disk_name*

   - Rekey all the disks at once:

     disk encrypt rekey 0x0 *

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

   disk encrypt show

Example

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

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

Shutting down a node running Data ONTAP operating in 7-Mode

When performing maintenance on a system running Data ONTAP operating in 7-Mode, you must shut down the node. Depending on your system's configuration, you might also need to turn off the power supplies.

About this task

Your system's configuration determines whether you turn off the power supplies after shutting down the node:

- 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 a stand-alone configuration, you must turn off the power supplies in the impaired node chassis.

Shutting down a node in an HA pair

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

Steps

1. Check the HA status of the impaired node from either node in the HA pair that is displaying the ONTAP prompt:

   cf status

2. Take the appropriate action based on the takeover status of the node.

If the impaired node...	Then...
Has been taken over by the healthy node and is halted	Go to the next step.
Has not been taken over by the healthy node and is running	Take over the impaired node from the prompt of the healthy node: `cf takeover`

3. Wait for two minutes after takeover of the impaired node to confirm that the takeover was completed successfully.
4. With the impaired node showing the `Waiting for giveback` message or halted, shut it down, depending on your configuration:

If the Service Processor (SP)...	Then...
Is configured	Log in to the SP, and then turn off the power: `system power off`
Is not configured	At the prompt of the impaired node, press Ctrl-C and respond **`Y`** to halt the node.

Shutting down a node in a stand-alone configuration

For a node that is not configured with a high-availability (HA) partner, you must perform a clean shutdown (verifying that all data has been written to disk) and disconnect the power supplies.

Steps

1. Shut down the node if it is not already shut down:

   `halt -t 0`

2. Shut down the power supplies, and then unplug both power cords from the source.

   The system is ready for maintenance.

Verifying the new controller module has no content in NVRAM

You must check that the new controller module has no content in NVRAM before completing the replacement.

Steps

1. Check the NVRAM LED on the controller module.

   You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The NVMEM LED is marked with a battery symbol and is located about 3" to the left of the label showing the MAC address on the controller module.

2. If the NVRAM LED is not flashing, there is no content in the NVRAM; You can skip the following steps and proceed to the next task in this procedure.
3. If the NVRAM LED is flashing, there is data in the NVRAM and you must disconnect the battery to clear the memory:

   a. If you are not already grounded, properly ground yourself.

   b. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and unplug the battery cable from the socket.

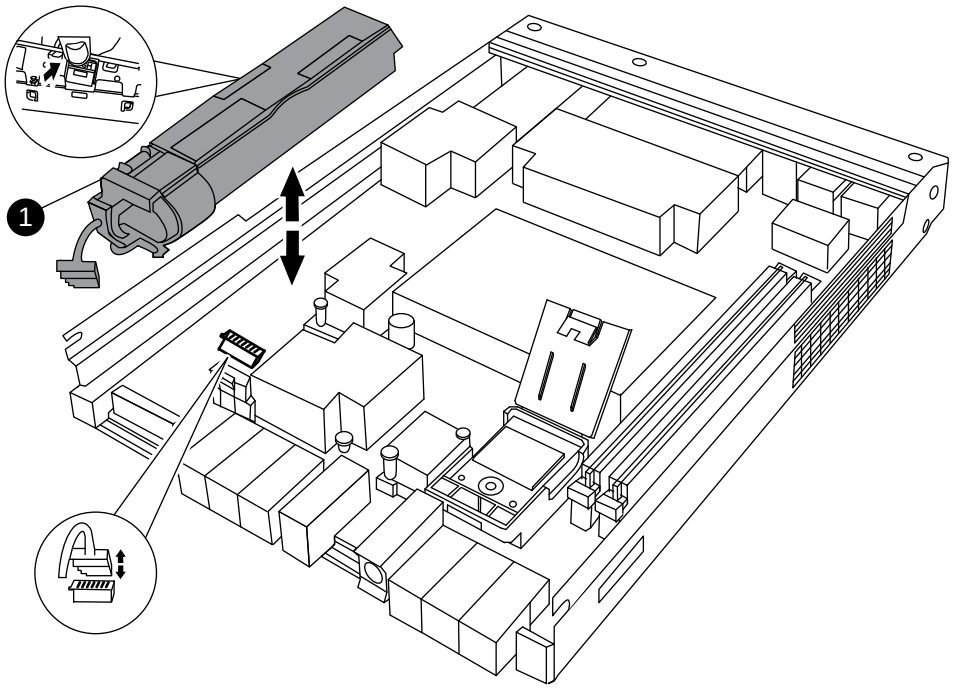

1	NVMEM battery

c. Confirm that the NVRAM LED is no longer lit.

d. Reconnect the battery connector.

4. Return to step 1 of this procedure to recheck the NVRAM LED.

Replacing the controller module hardware

To replace the controller module hardware, you must remove the impaired node, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

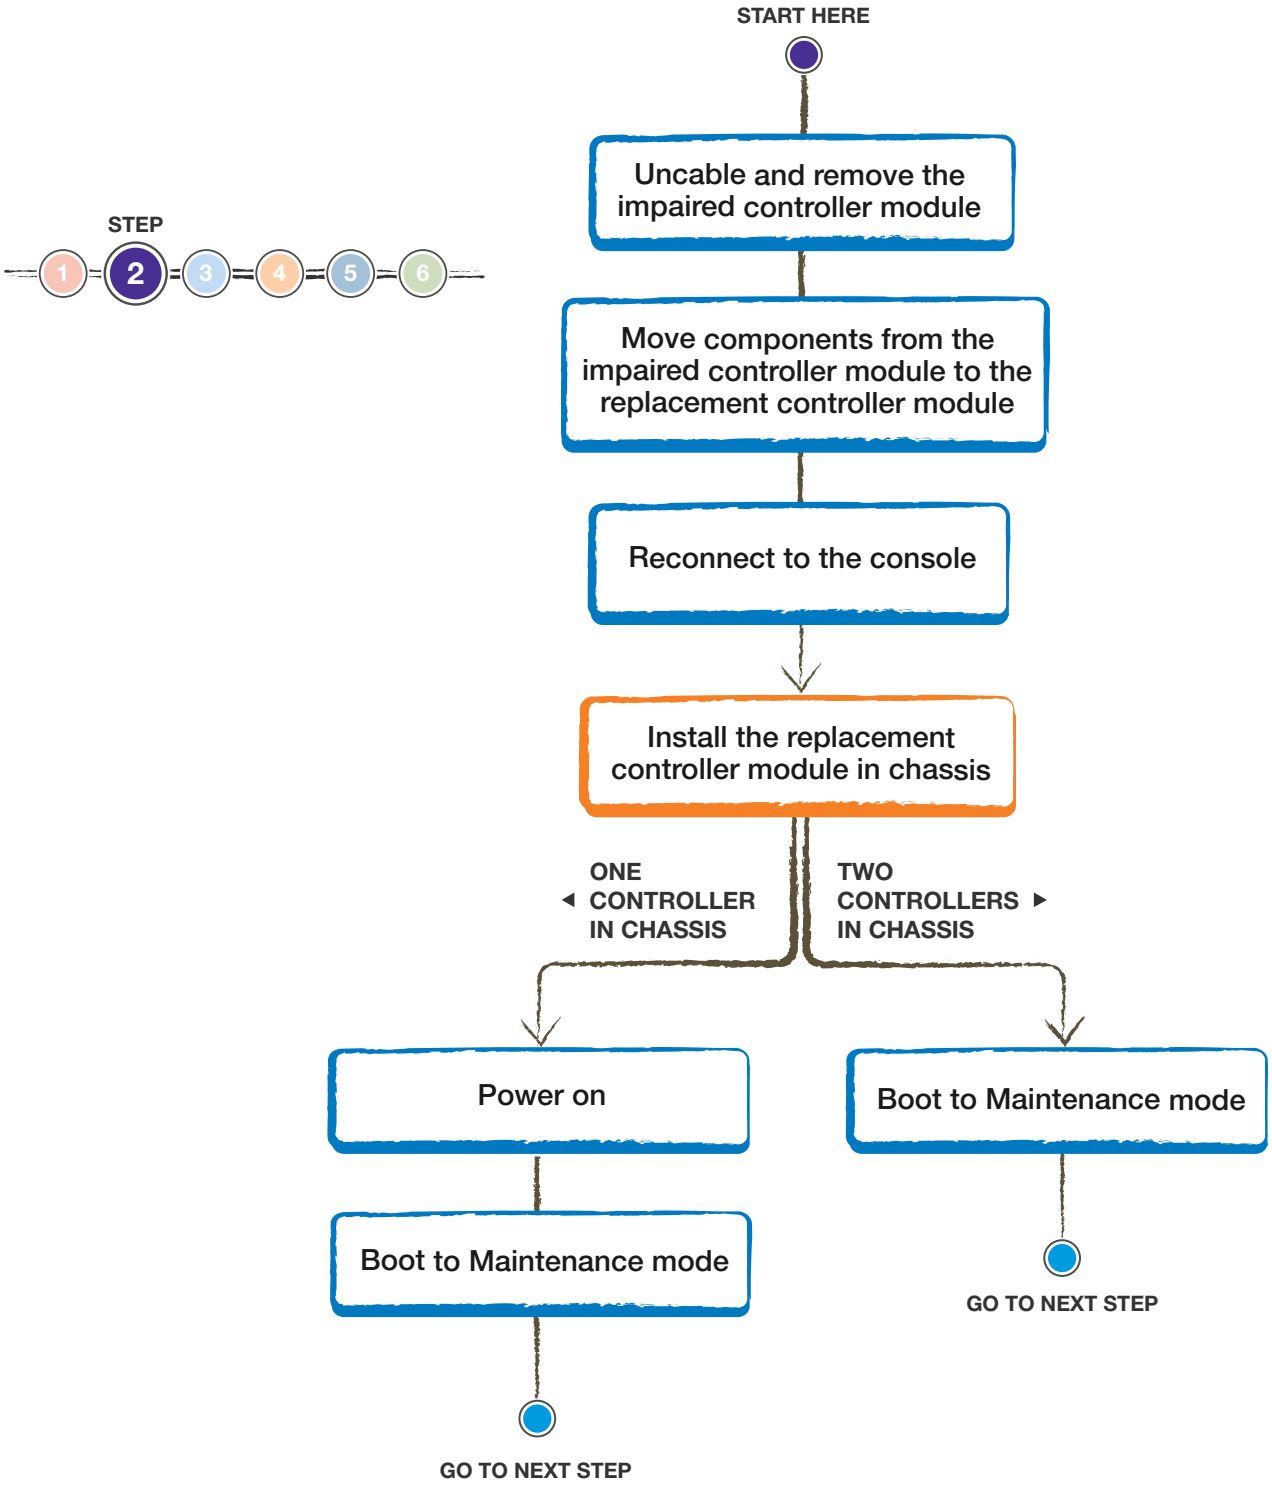

Steps

1. Removing the controller module and moving the components on page 9
2. Installing the new controller module and booting the system on page 14

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 for two minutes after takeover of the impaired node to confirm 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.

Steps

1. Removing the controller module from the system on page 9
2. Moving the boot device on page 11
3. Moving the NVRAM battery on page 12
4. Moving the DIMMs to the new controller module on page 12

Removing the controller module from the system

To replace the controller module, you must first remove the old controller module from 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 arm, 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 arm so that when you reinstall the cable management arm, the cables are organized.
3. Remove the cable management arms from the left and right sides of the controller module and set them aside.

   The illustration shows the cable management arms on a FAS2552 system. The procedure is the same for all FAS25xx systems.

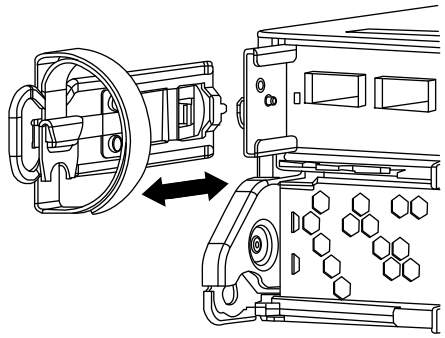

4. Squeeze the latch on the cam handle until it releases, as shown in the following illustration. Open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

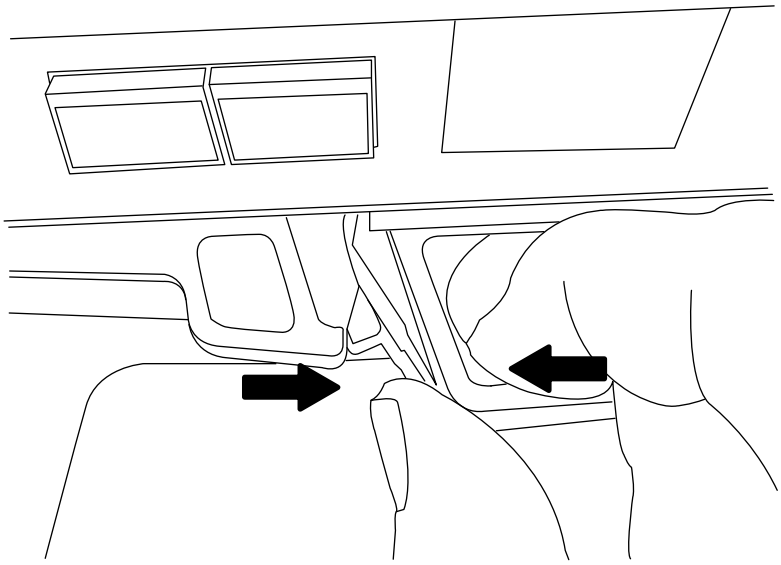

5. Turn the controller module over and open it by pressing the button to release the cover, and then slide the cover out.

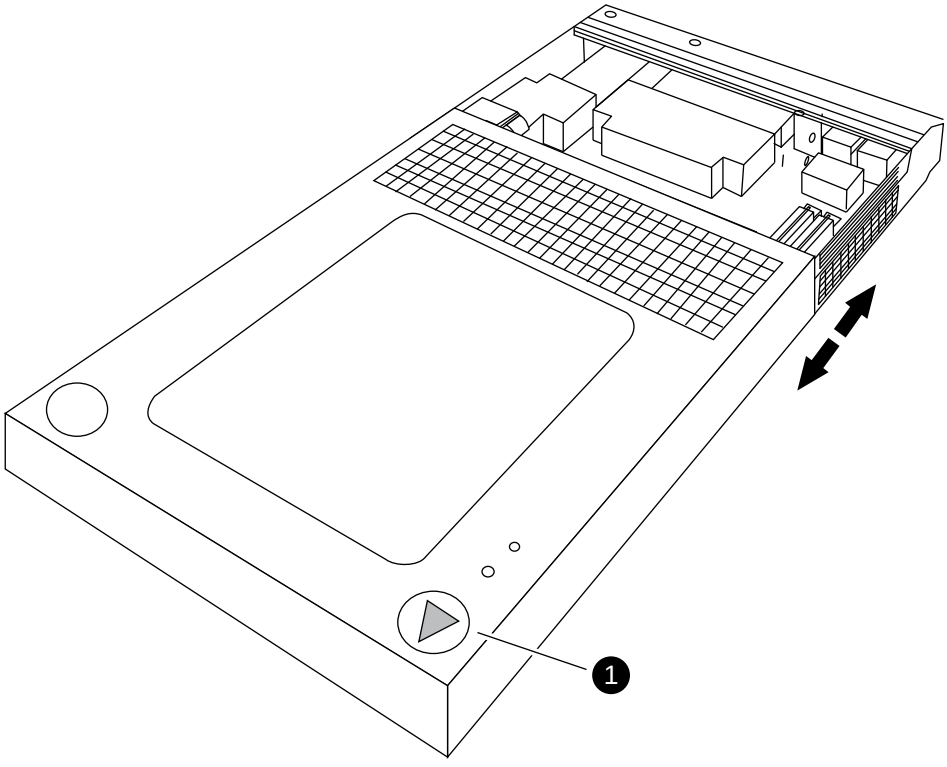

1	Button to release controller module cover

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:

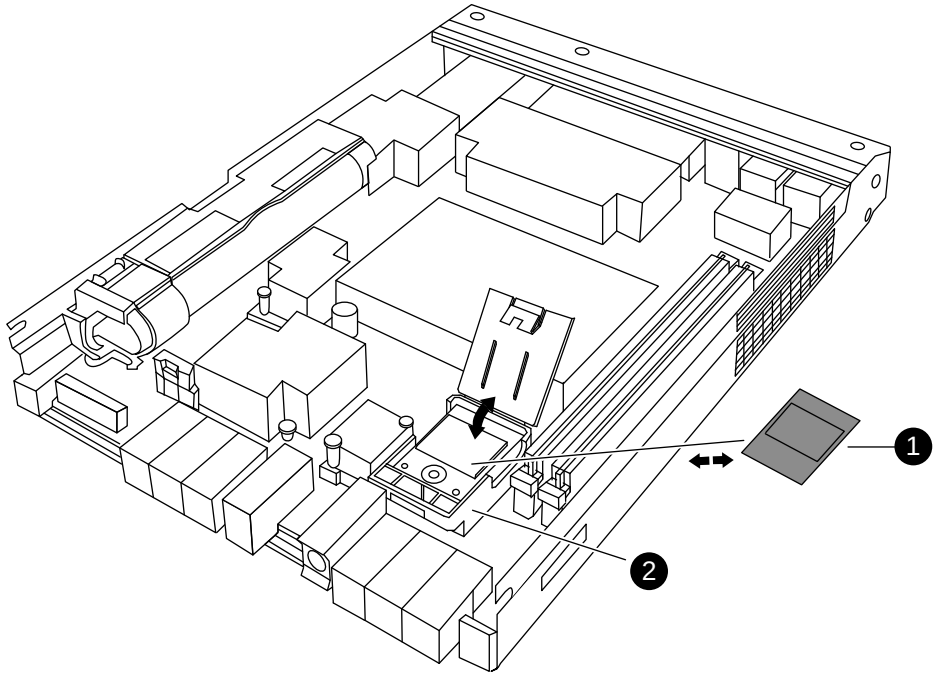

1	Boot device
2	Boot device holder; not removable

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, 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. Verify that the boot device is seated squarely and completely in the socket or connector.

   If necessary, remove the boot device and reseat it into the socket.

6. Close the boot device cover.

Moving the NVRAM battery

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

Steps

1. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

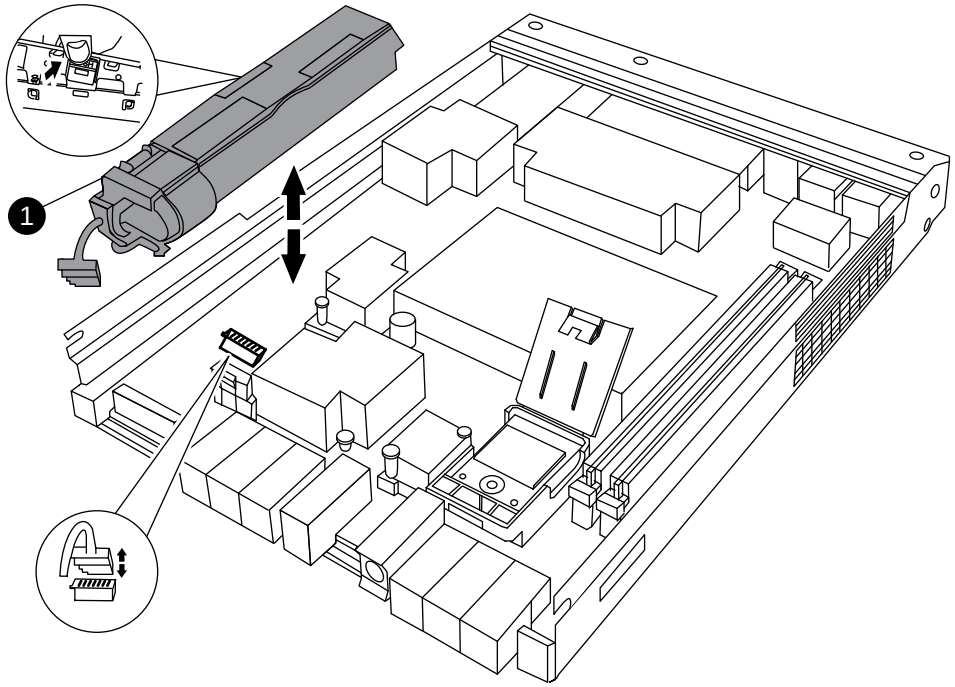

1	NVMEM battery

2. Grasp the battery and press the tab marked PUSH, and then lift the battery out of the holder and controller module.
3. In the new controller module, seat the battery in the holder.

   Attention: Do not connect the NVRAM keyed battery plug into the socket until after the NVRAM DIMM has been installed.

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. Verify that the NVMEM battery cable connector is not plugged into the socket.
2. Locate the DIMMs.

   If you are moving DIMMs on a FAS25xx system:

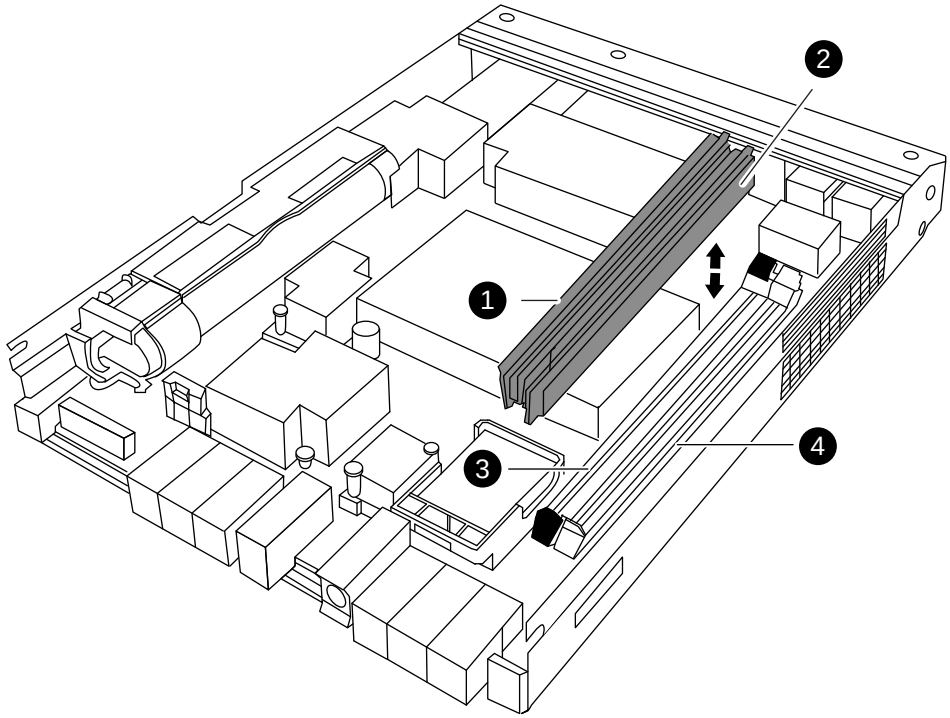

1	System DIMM
2	NVMEM DIMM The NVMEM DIMM has an *NVMEM* label on one of the chips.
3	System DIMM slot
4	NVMEM DIMM slot The NVMEM DIMM slot has white ejector tabs.

3. 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.
4. 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.

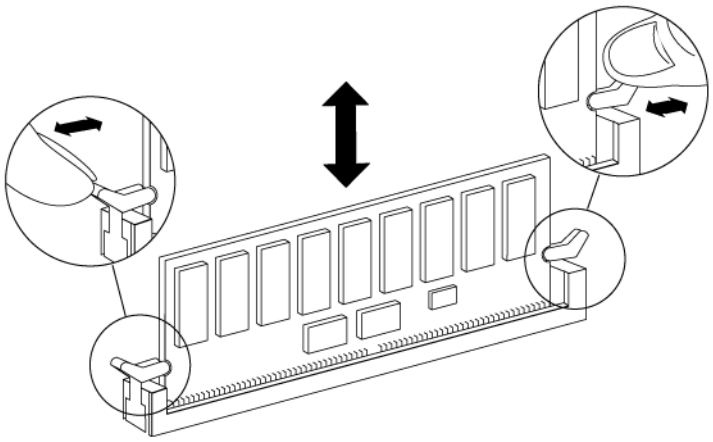

Attention: You must carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Locate the corresponding slot for the DIMM in the new controller module, align the DIMM over the slot, and then 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, you should realign the DIMM with the slot and reinsert it.

   Important: You must install the NVMEM DIMM only in the NVMEM DIMM slot.

6. 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.

7. 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.

8. Repeat these steps to move additional DIMMs, as required.

9. In the new controller module, orient the NVMEM battery cable connector to the socket on the controller module and plug the cable into the socket.

   You must ensure that the plug locks down onto the socket on the controller module.

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 might update the system firmware when it boots. You must 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: You must not completely insert the controller module in the chassis until instructed to do so.

2. Recable the management port or serial console port so that 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	**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:** You must not use excessive force when sliding the controller module into the chassis; you might damage the connectors. **b.** Enter one of the following commands from the healthy node's console and wait for the giveback to complete: For systems operating in... / Issue the command... 7-Mode / `cf giveback` Clustered Data ONTAP / • In Data ONTAP 8.2 or later: `storage failover giveback –ofnode impaired_node_name` **c.** If you have not already done so, reinstall the cable management arm, and then tighten the thumbscrew on the cam handle on back of the controller module. **d.** Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	**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:** You must 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. **c.** If you have not already done so, reinstall the cable management arm, and then tighten the thumbscrew on the cam handle on back of the controller module. **d.** 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 an HA configuration you must confirm that the healthy node remains down.

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

Restoring and verifying the system configuration after hardware replacement

After replacing the hardware components, you should verify the low-level system configuration of the replacement controller and reconfigure FC settings if necessary.

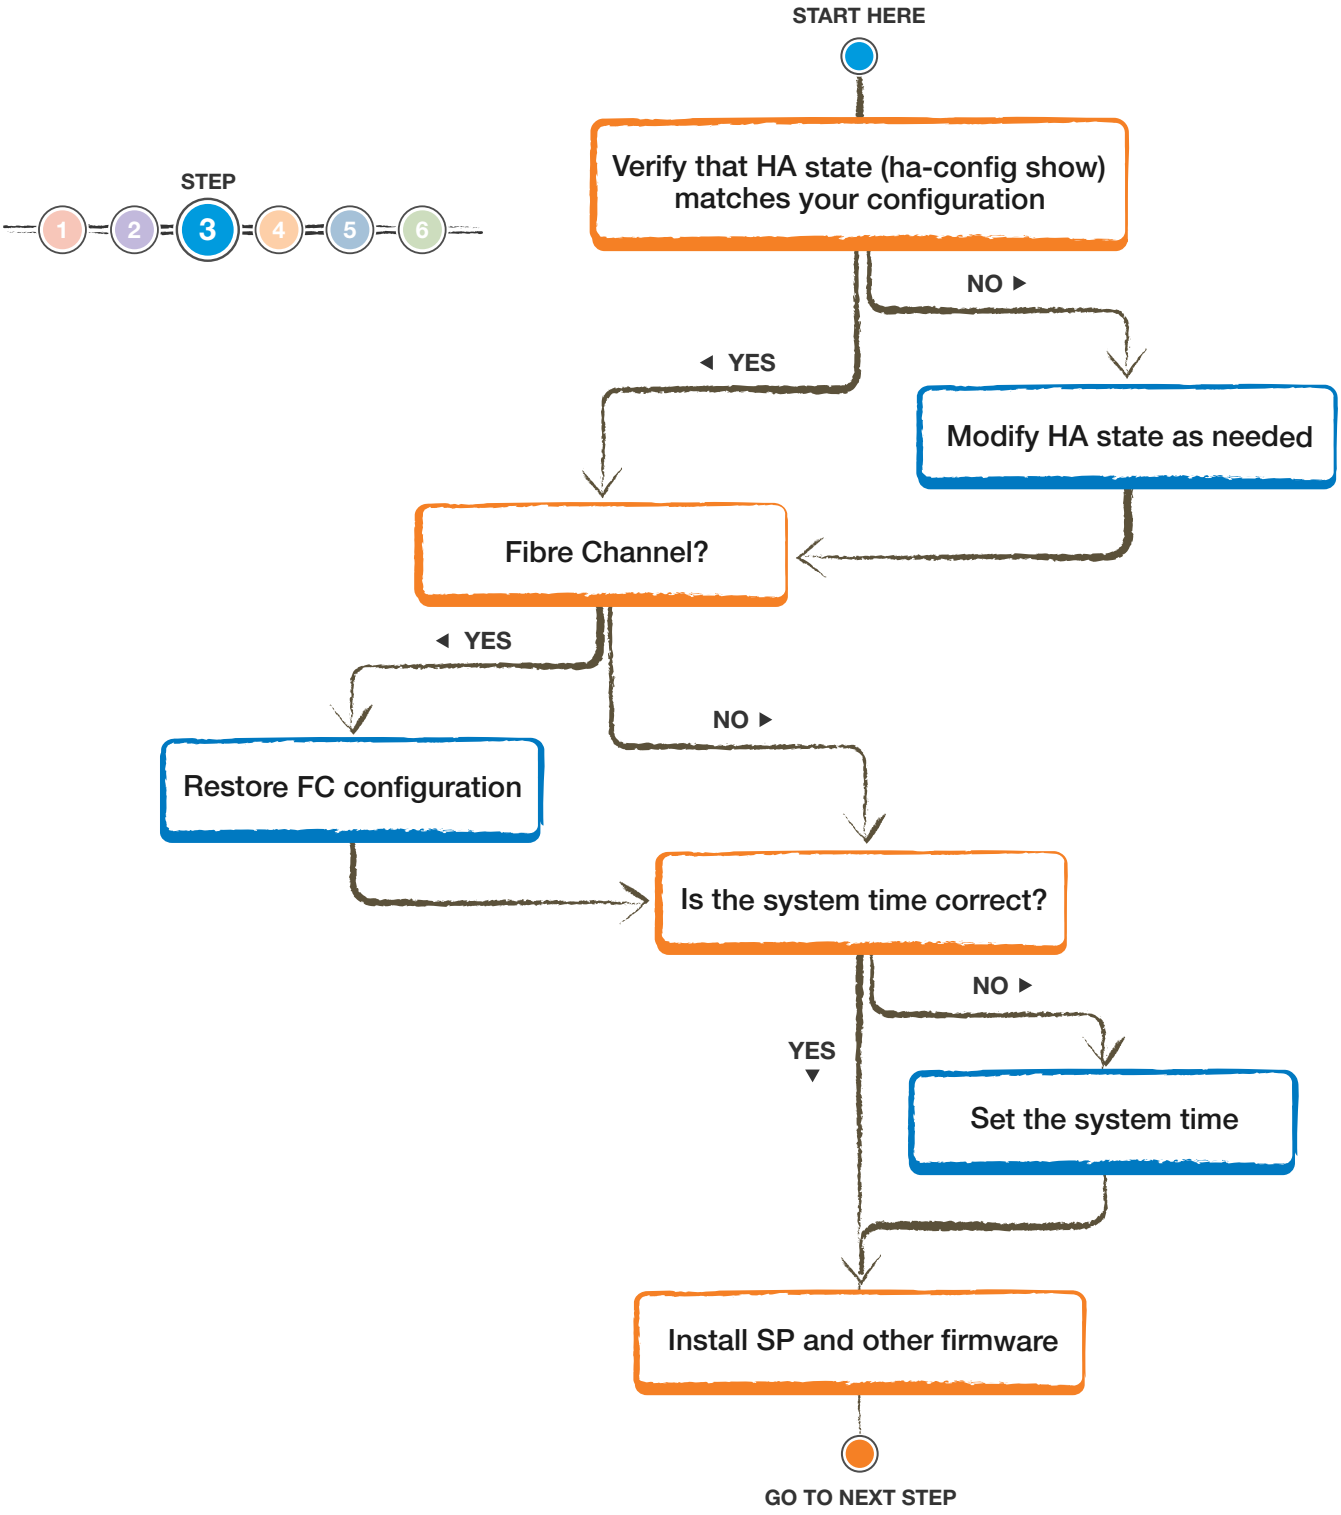

Steps

1. Verifying and setting the HA state of the controller module on page 17
2. Restoring Fibre Channel configurations (7-Mode CNA) on page 17
3. Setting the system time after replacing the controller module on page 18
4. Installing the firmware after replacing the controller module on page 19

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, 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 | non-ha]

If your system is...	Run 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 | non-ha]

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

Restoring Fibre Channel configurations (7-Mode CNA)

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

Before you begin

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

Steps

1. If you have not already done so, reboot the replacement node to Maintenance mode by pressing **Ctrl-C** when you see the message `Press Ctrl-C for Boot Menu`.

2. Answer **y** when prompted by the system.
3. Select the Maintenance mode option from the menu.
4. Restore the FC port configuration:

   ucadmin modify -t *initiator|target adapter_name*

   - Enter *initiator* if you are connecting to a Fibre Channel tape device.
   - Enter *target* if you are in a SAN configuration.
5. Take one of the following actions, depending on your configuration:

If the FC port configuration is...	Then...
The same for both ports	Answer **y** when prompted by the system. Modifying one port in a port pair modifies the other port as well.
Different	a. Answer **n** when prompted by the system. b. Restore the FC port configuration: **ucadmin modify -t *initiator|target adapter_name***

6. Exit Maintenance mode:

   halt

   After you issue the command, you must wait until the system stops at the LOADER prompt.
7. Boot the replacement node:

   boot_ontap
8. Verify the values of the variables:

   ucadmin show

Setting the system time after replacing the controller module

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, verify 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. Determine the system time by using the `date` command on the healthy node (if the system is in an HA pair) or another reliable time source.
3. Set the date in GMT on the replacement node:

   set date *mm/dd/yyyy*

4. Set the time in GMT on the replacement node:

   `set time hh:mm:ss`

Installing the firmware after replacing the controller module

After replacing the controller module, you must install the latest firmware if your system is running a version of Data ONTAP earlier than 8.2, 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

If your system is running ONTAP 8.2 or later, the SP firmware and BIOS automatically update to the baseline image included with the ONTAP version. Other system firmware from the old controller module still resides on the boot device and typically does not need updating.

If your system is running ONTAP 8.2 or later, you should skip this procedure.

Steps

1. Check the configuration of the SP from the LOADER prompt:

   `sp status`

   For the latest release of SP firmware, log in to the NetApp Support Site at *mysupport.netapp.com* and update it, if needed, in the following steps.

2. Log in to the SP from an administration host:

   `ssh username@SP_IP_address`

3. Download and install the most current version of firmware for your system by 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.

Related information

Find a System Administration Guide for your version of ONTAP 9
Find a System Administration Guide for your version of Data ONTAP 8

Running diagnostics tests after replacing a controller module

You should run focused diagnostic tests for specific components and subsystems whenever you replace a component of the controller.

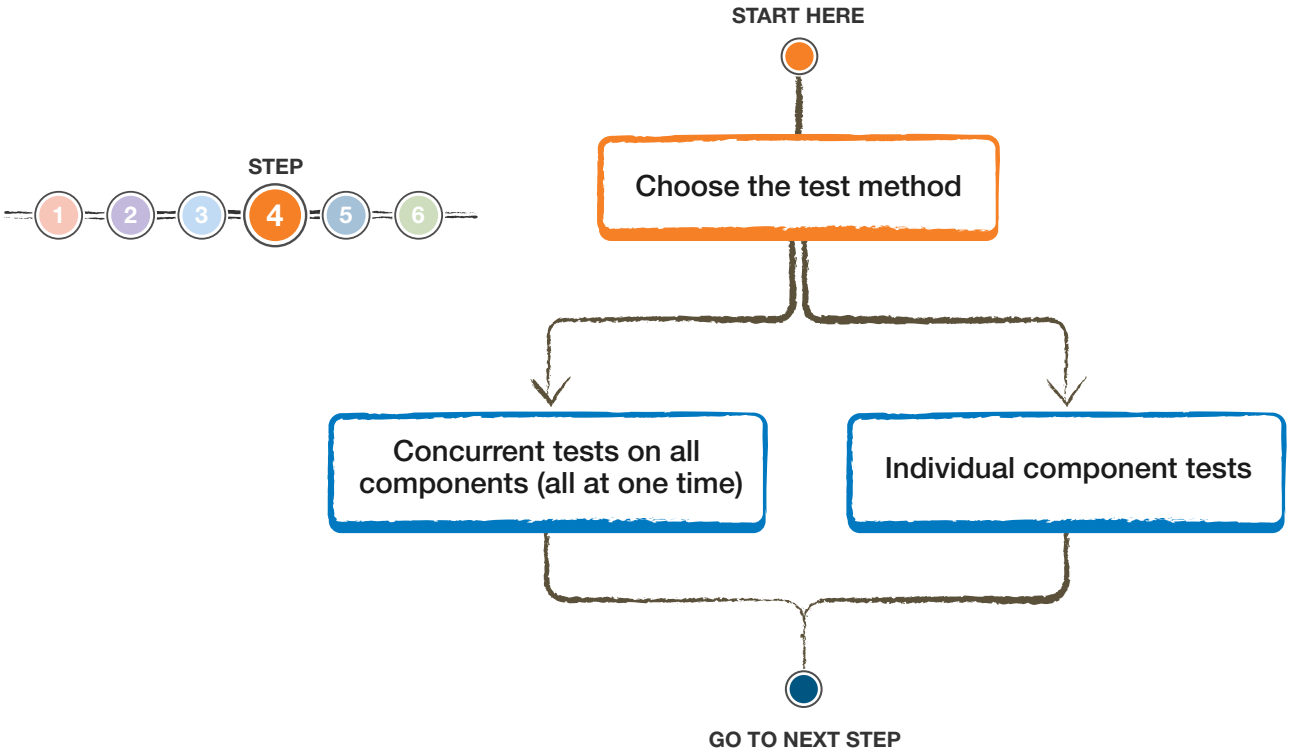

Before you begin

- Your system must be at the LOADER prompt to start system-level diagnostics.
- For ONTAP 8.2 and later, you do not require 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, run the system-level diagnostic test: **`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 that are designed specifically for system-level diagnostics.

   Important: During the `boot_diags` process, you might see a prompt warning that when entering Maintenance mode in an HA configuration, you must confirm that the partner remains down. To continue to Maintenance mode, you should enter **`y`**
3. Clear the status logs: **`sldiag device clearstatus`**
4. Display and note the available devices on the controller module: **`sldiag device show -dev mb`**

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

   - **`bootmedia`** is the system booting device.

- **cna** is a Converged Network Adapter or interface that is not connected to a network or storage device.
- **env** is the motherboard environmentals.
- **mem** is the system memory.
- **nic** is a network interface card.
- **nvmem** is a hybrid of NVRAM and system memory.
- **sas** is a Serial Attached SCSI device that is not connected to a disk shelf.

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

Choices

- Running diagnostics tests concurrently after replacing the controller module on page 21
- Running diagnostics tests individually after replacing the controller module on page 22

Running diagnostics tests concurrently after replacing the controller module

After replacing the controller module, you can 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 that you make. If you run more tests in addition to the default tests, the diagnostic test process takes longer to complete.

Steps

1. Display and note the available devices on the controller module: **sldiag device show -dev mb**

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

   - **bootmedia** is the system booting device.
   - **cna** is a Converged Network Adapter or interface that is not connected to a network or storage device.
   - **env** is the motherboard environmentals.
   - **mem** is the system memory.
   - **nic** is a network interface card.
   - **nvmem** is a hybrid of NVRAM and system memory.
   - **sas** is a Serial Attached SCSI device that is not connected to a disk shelf.

2. Review the enabled and disabled devices in the output from step 1 and then determine which tests you want to run concurrently.
3. List the individual tests for each device:

   sldiag device show -dev *dev_name*

4. Verify that the tests were modified: **sldiag device show**
5. Repeat steps 2 through 4 of this procedure for each device.
6. Run diagnostics on all the devices: **sldiag device run**

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

The tests are complete when the following message is displayed:

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

7. After the tests are complete, verify that there are no hardware problems on your storage system:

   sldiag device status -long -state failed

8. Correct any issues that are found, and repeat this procedure.

Running diagnostics tests individually after replacing the controller module

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

Steps

1. Clear the status logs: **sldiag device clearstatus**

2. Display the available tests for the selected devices:

Device type	Command
boot media	**sldiag device show -dev bootmedia**
cnafcal	**sldiag device show -dev cna** **sldiag device show -dev fcal**
env	**sldiag device show -dev env**
mem	**sldiag device show -dev mem**
nic	**sldiag device show -dev nic**
nvmem	**sldiag device show -dev nvmem**
sas	**sldiag device show -dev sas**

3. 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.

4. 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

5. Run the selected tests:

Device type	Command
boot media	**sldiag device run -dev bootmedia**
cnafcal	**sldiag device run -dev cna** **sldiag device run -dev fcal**
env	**sldiag device run -dev env**
mem	**sldiag device run -dev mem**

Device type	Command
nic	**sldiag device run -dev nic**
nvmem	**sldiag device run -dev nvmem**
sas	**sldiag device run -dev sas**

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

```
<SLDIAG:_ALL_TESTS_COMPLETED>
```

6. Verify that no tests failed:

Device type	Command
boot media	**sldiag device status -dev bootmedia -long -state failed**
cnafcal	**sldiag device status -dev cna** **sldiag device status -dev fcal**
env	**sldiag device status -dev env -long -state failed**
mem	**sldiag device status -dev mem -long -state failed**
nic	**sldiag device status -dev nic -long -state failed**
nvmem	**sldiag device status -dev nvmem**
sas	**sldiag device status -dev sas -long -state failed**

Any tests that failed are displayed.

7. 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: **sldiag device clearstatus** **b.** Verify that the log is cleared: **sldiag device status** The following `SLDIAG: No log messages are present.` default response is displayed. You have completed system-level diagnostics.

If the system-level diagnostics tests...	Then...
Resulted in some test failures	Determine the cause of the problem: **a.** Exit Maintenance mode: **`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 then unplug them from the power sources. **c.** Check the controller module you are servicing to verify that you have observed all of 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 then turn them on. **e.** Select **Boot to Maintenance mode** from the menu. **f.** Exit Maintenance mode: **`halt`** After you issue the command, you must wait until the system stops at the LOADER prompt. **g.** Enter **`boot_diags`** at the prompt, and then rerun the system-level diagnostic test.

8. Exit system-level diagnostics, and continue with recabling and restoration of the storage system.

Completing the recabling and final restoration of operations

To complete the replacement procedure, you must recable the storage system, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller.

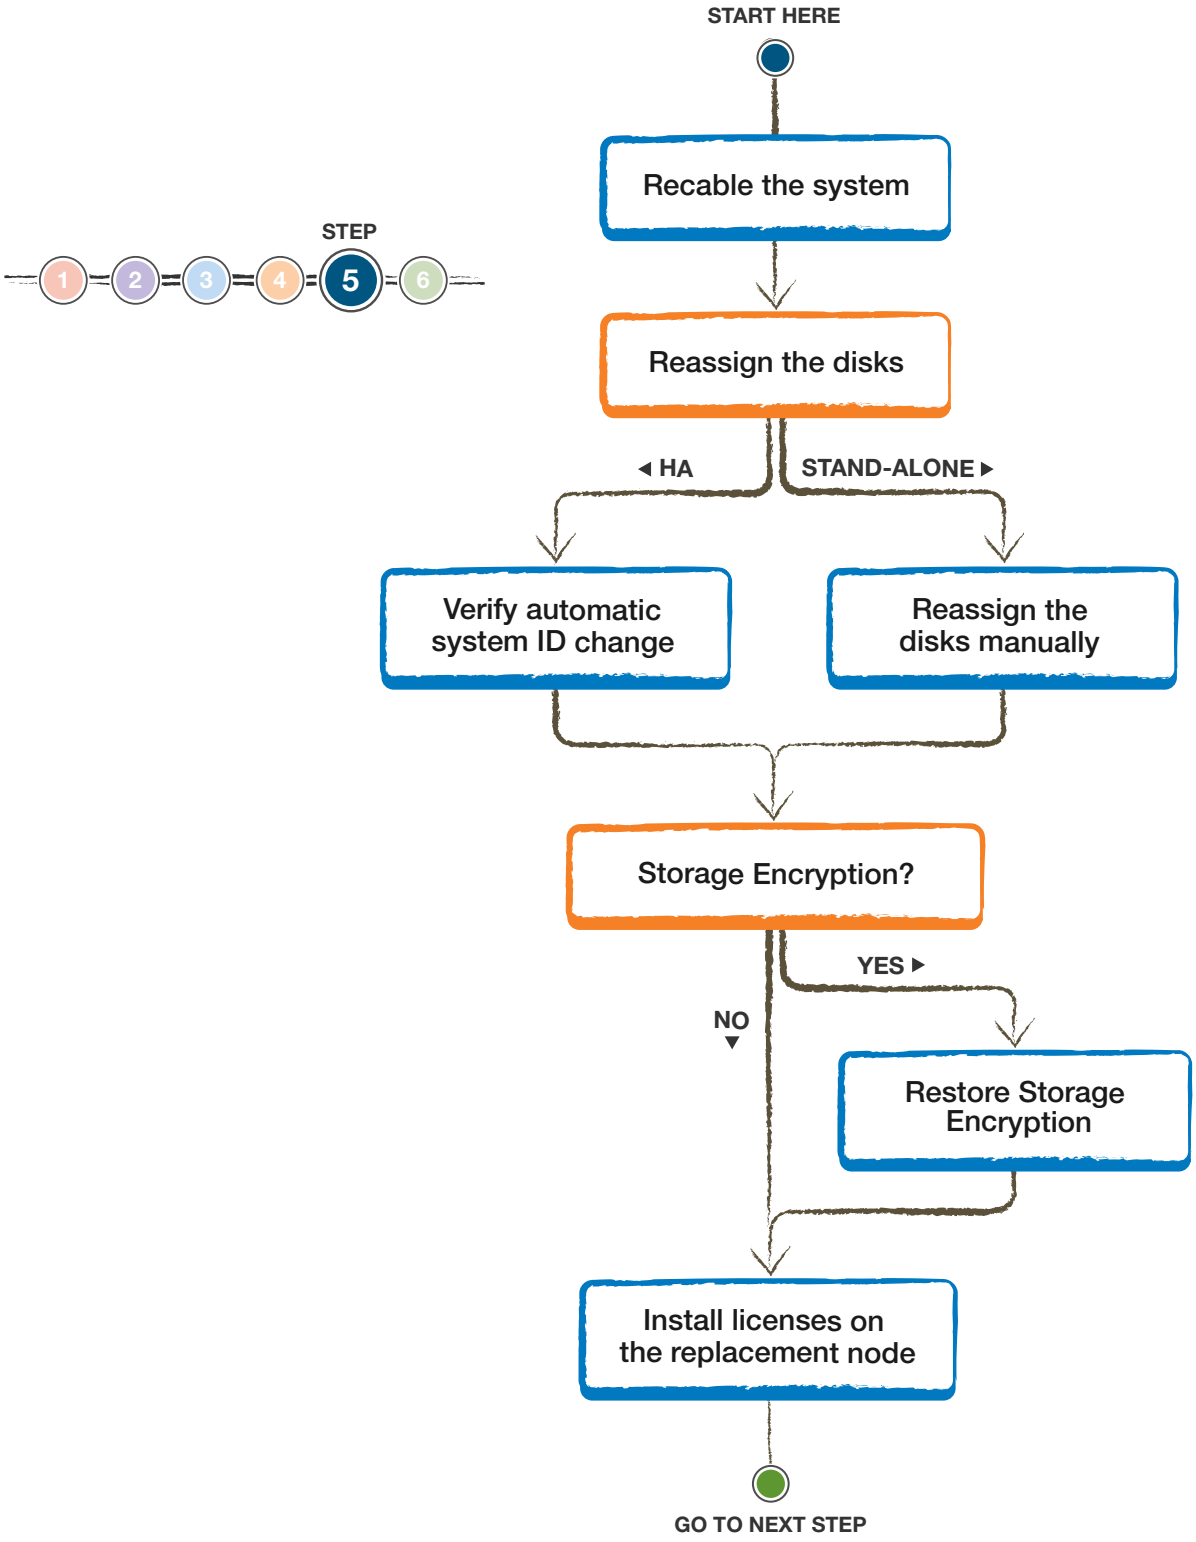
START HERE
STEP
1
2
3
4
5
6
Recable the system
Reassign the disks
◄ HA
STAND-ALONE ►
Verify automatic system ID change
Reassign the disks manually
Storage Encryption?
YES ►
NO
▼
Restore Storage Encryption
Install licenses on the replacement node
GO TO NEXT STEP

Steps

1. Recabling the system on page 26
2. Reassigning disks on page 26
3. Restoring Storage Encryption functionality after replacing controller modules on page 29
4. Installing licenses for replacement nodes operating in 7-Mode on page 29

Recabling the system

After running diagnostics, you must recable the storage and network connections of the controller module.

Steps

1. Reinstall the cable management arms and recable the controller module, 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:

      ToolChest

      The “Quick Start Guide” provides instruction to collect and analyze data from your system.

      Config Advisor Quick Start Guide

   b. Check the rules for “SAS Cabling Checks” and then examine the output from Config Advisor.

      You must verify that all of the disk shelves are displayed and that all disks appear in the output. You must correct any cabling issues that you might find.

Related information

Disk Shelves Documentation

Reassigning disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

About this task

You must use the correct procedure for your configuration:

If the controller is in...	Then use this procedure...
An HA pair	*Verifying the system ID change on a system operating in 7-Mode* on page 26
A stand-alone configuration	*Manually reassigning the system ID on a stand-alone system in 7-mode* on page 28

Verifying the system ID change on an HA system operating in 7-Mode

You must confirm the system ID change when you boot the replacement node, and then verify that the change was implemented.

About this task

This procedure applies only to systems that are in an HA pair and are running Data ONTAP operating in 7-Mode.

Steps

1. If the replacement node is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode:

   `halt`

   After you issue the command, you must wait until the system stops at the LOADER prompt.

2. From the LOADER prompt on the replacement node, display the Boot menu:

   a. Boot the replacement node:

      `boot_ontap`

   b. Press **`Ctrl-c`** when prompted to display the Boot menu.

3. Wait until the `Waiting for giveback...` message is displayed on the console of the replacement node and then, on the healthy node, verify that the controller module replacement has been detected and that the new partner system ID has been automatically assigned:

   `cf status`

   You should see a message similar to the following, which indicates that the system ID change has been detected:

   ```
   HA mode.
   System ID changed on partner (Old: 1873774576, New: 1873774574).
   partner_node has taken over target_node.
   target_node is ready for giveback.
   ```

   The message shows the new system ID of the replacement node. In this example, the new system ID is 1873774574.

4. From the healthy node, verify that all coredumps are saved: **`partner savecore`**

   If the command output indicates that a savecore is in progress, you must wait for the savecore to finish before initiating the giveback operation. You can monitor the progress of the savecore: **`partner savecore -s`**

5. Initiate the giveback operation after the replacement node displays the `Waiting for Giveback...` message:

   `cf giveback`

   You should see a message similar to the following noting the system ID change and prompting you to continue:

   ```
   System ID changed on partner. Giveback will update the ownership of partner disks with
   system ID: 1873774574.
   Do you wish to continue {y|n}?
   ```

   You must enter **`y`** to proceed. If the giveback is vetoed, you can consider overriding the veto.

   ONTAP 9 High-Availability Configuration Guide

   Data ONTAP 8.2 High Availability and MetroCluster Configuration Guide for 7-Mode

6. Verify that the disks were assigned correctly:

   `disk show`

   You must verify that the disks belonging to the replacement node show the new system ID for the replacement node. In the following example, the disks owned by node2 now show the new system ID, 1873774574:

   #### Example

   ```
   system-1> disk show
     DISK            OWNER         POOL   SERIAL NUMBER  HOME           DR HOME
   ---------------   ------------- -----  -------------  -------------  ------------------
   disk_name node2    (1873774574) Pool0  J8Y0TDZC       system-2       (1873774574)
   ```

```
disk_name node1     (118065578) Pool0  J8Y09DXC       system-1        (118065578)
.
.
.
```

7. Verify that the expected volumes are present and are online for each node:

   vol status

Manually reassigning the system ID on a stand-alone system operating in 7-Mode

In a stand-alone system, you must manually reassign disks to the new controller's system ID before you return the system to normal operating condition.

About this task

This procedure applies to stand-alone systems that are operating in 7-Mode.

Steps

1. If you have not already done so, reboot the replacement node, interrupt the boot process by entering **Ctrl-C**, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter **Y** when prompted to override the system ID due to a system ID mismatch.

2. View the system IDs:

   disk show -a

   Note: Make a note of the old system ID, which is displayed as part of the disk owner column.

   #### Example

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

 DISK     OWNER        POOL    SERIAL NUMBER   HOME
--------  -----------  ------  -------------   -------------
system-1  (118073209)  Pool0   J8XJE9LC        system-1  (118073209)
system-1  (118073209)  Pool0   J8Y478RC        system-1  (118073209)
.
.
.
```

3. Reassign disk ownership by using the system ID information obtained from the `disk show` command:

   disk reassign -s *old system ID*

   In the case of the preceding example, the command is **disk reassign -s 118073209**.

   You can respond **Y** when prompted to continue.

4. Verify that the disks were assigned correctly: **disk show -a**

   You must verify that the disks belonging to the replacement node show the new system ID for the replacement node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

   #### Example

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER          POOL   SERIAL NUMBER  HOME
```

```
-------    -------------  -----  -------------  -------------
system-1  (118065481)     Pool0  J8Y0TDZC       system-1  (118065481)
system-1  (118065481)     Pool0  J8Y09DXC       system-1  (118065481)
.
.
.
```

5. If the replacement node is in Maintenance mode (showing the *> prompt), exit Maintenance mode:

   halt

   After you issue the command, you must wait until the system stops at the LOADER prompt.

6. Boot the operating system: **boot_ontap**

Restoring Storage Encryption functionality after replacing controller modules

After replacing the controller module for a storage system that you previously configured to use Storage Encryption, you must perform additional steps to restore Storage Encryption functionality in an uninterrupted way. You can skip this task on storage systems that do not have Storage Encryption enabled.

Steps

1. Reconfigure Storage Encryption at the storage system prompt: **key_manager setup**
2. Complete the steps in the setup wizard to configure Storage Encryption.

   You must verify that a new passphrase is generated, and you must select **Yes** to lock all drives.

3. Repeat step *1* on page 29 and step *2* on page 29 on the partner node.

   You should not proceed to the next step until you have completed the Storage Encryption setup wizard on each node.

4. On each node, verify that all disks are rekeyed: **disk encrypt show**

   None of the disks should list a key ID of 0x0.

5. On each node, load all authentication keys: **key_manager restore -all**
6. On each node, verify that all keys are stored on their key management servers: **key_manager query**

   None of the key IDs should have an asterisk next to it.

Installing licenses for replacement nodes operating in 7-Mode

You must reinstall new license keys for replacement nodes 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 requires 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 knowledgebase article 3013749 at *NetApp Knowledgebase Answer 1002749: Data ONTAP 8.2 and 8.3 Licensing Overview and References* and the *Data ONTAP System Administration Guide for 7-Mode*.

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

You have a 90-day grace period 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 require 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, you should contact technical support.

2. You must wait until the ONTAP command-line interface has been up for at least five minutes and then confirm that the license database is running.
3. Install the license keys:

   `license add license_key license_key license_key...`

   You can add one license or multiple licenses simultaneously, with each license key separated by a comma or a space.

   If the ONTAP command-line interface was not up for a sufficient amount of time, you might receive a message indicating that the license database is unavailable.
4. Verify that the licenses have been installed:

   `license show`

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.

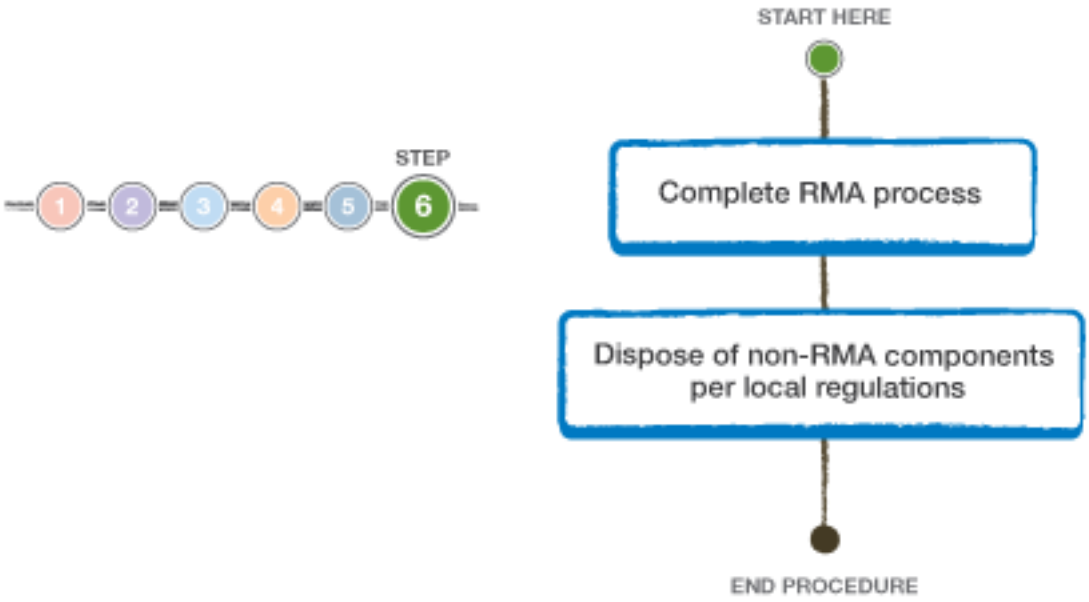

Related information

NetApp Support

Disposing of batteries

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

Related information

https://library.netapp.com/ecm/ecm_download_file/ECMP12475945

Replacing a controller module in ONTAP

You must follow a specific series of steps to replace the depending on your mode and version of ONTAP.

Steps

1. Preparing the system for controller replacement on page 31
2. Replacing the controller module hardware on page 39
3. Restoring and verifying the system configuration after hardware replacement on page 47
4. Running diagnostics tests after replacing a controller module on page 51
5. Completing the recabling and final restoration of operations on page 55
6. Completing the replacement process on page 63

Preparing the system for controller replacement

You must gather information as shown in the workflow diagram, do a take-over and then proceed to shut down the impaired node in an HA pair.

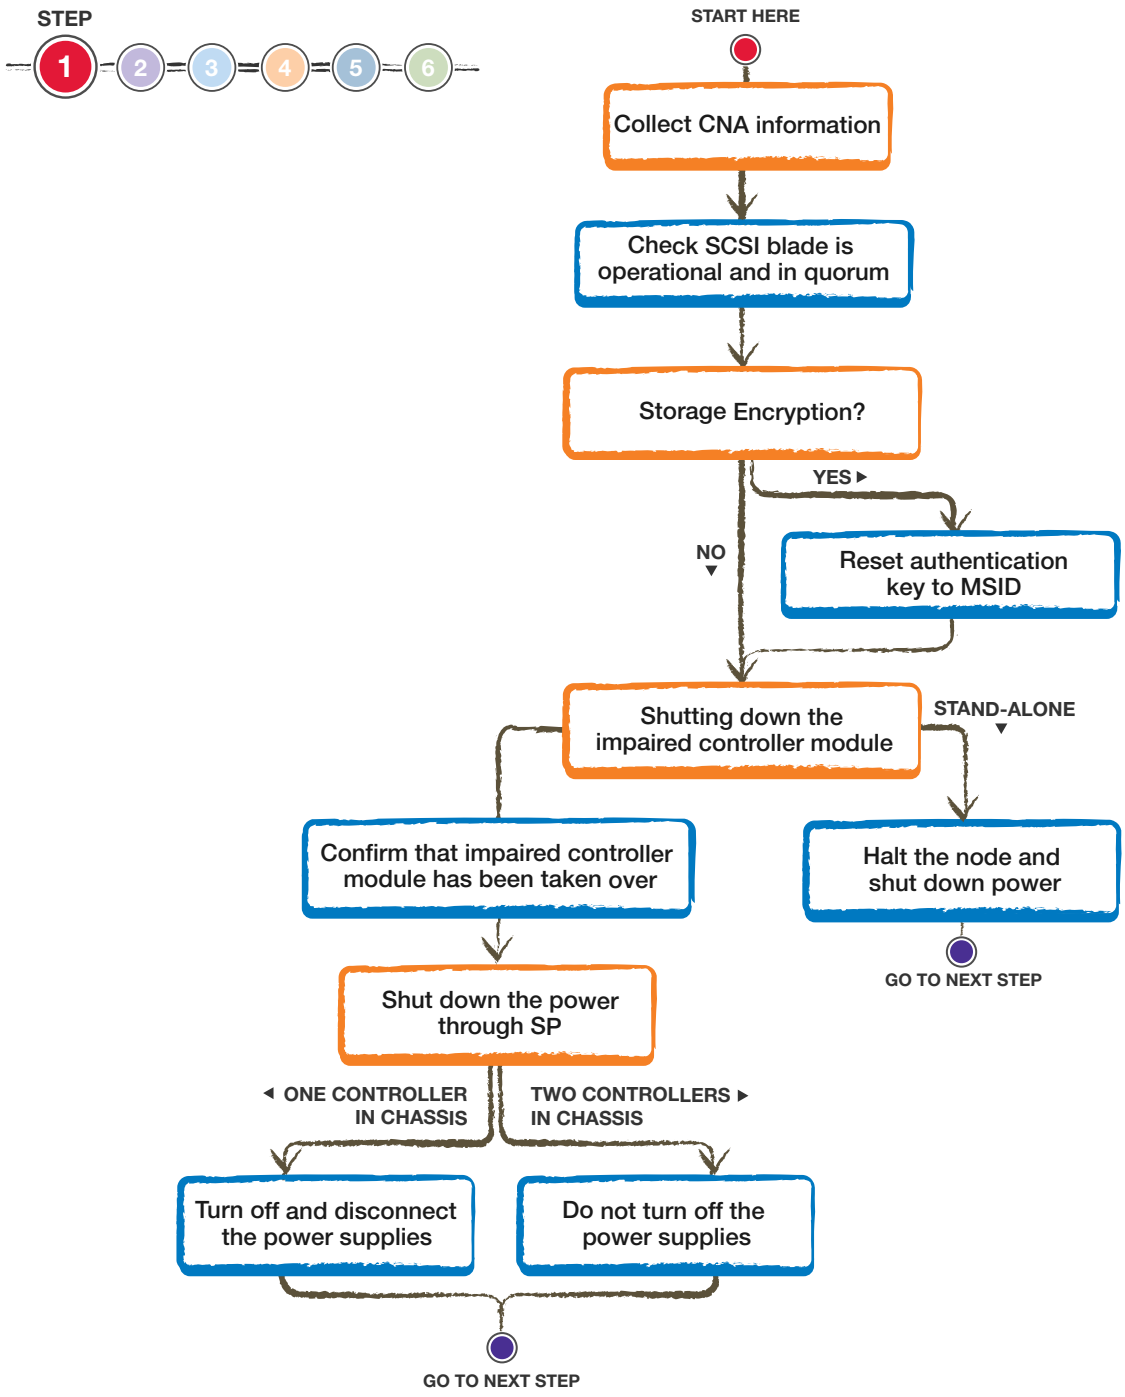

Steps

1. Determining your controller CNA port configuration on page 33
2. Checking quorum on the SCSI blade on page 33
3. Preparing for Storage or Volume Encryption configurations on page 34
4. Shutting down the target controller on page 35
5. Verifying the new controller module has no content in NVRAM on page 37

Determining your controller CNA port configuration

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 the information on the new controller module. You must also check whether the SCSI process is in quorum with the other nodes in the cluster.

Steps

1. Save the port configuration information for the impaired node:

If your system is running...	Then...
Data ONTAP 8.2.1 and earlier	Run the following command on the console of the healthy node: **`partner fcadmin config`**
ONTAP 8.2.2 and later	**a.** Run the following command on the console of the impaired node: **`system node hardware unified-connect show`** **b.** Run the following `Cluster-Mode` command on the console of the impaired node: **`system node hardware unified-connect modify`**

2. Copy and save the information displayed on the screen to a safe location for later reuse.

Checking quorum on the SCSI blade

Before you replace your controller module in an HA pair, you must check that the SCSI process is in quorum with other controller modules in the cluster.

Steps

1. Verify that 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
------------------- ---------------- ------------ --------------------------
11/1/2013 14:03:51 node1 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this node
established quorum with the other nodes in the cluster.
11/1/2013 14:03:51 node2 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this node
established quorum with the other nodes in the cluster.
11/1/2013 14:03:48 node3 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this node
established quorum with the other nodes in the cluster.
11/1/2013 14:03:43 node4 INFORMATIONAL scsiblade.in.quorum: The scsi-blade on this node
established quorum with the other nodes in the cluster.
```

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

Preparing for Storage or Volume Encryption configurations

If the storage system whose controller you are replacing is configured to use Storage or Volume Encryption, you must first reset the authentication keys of the disks to an MSID key (the default security ID set by the manufacturer). This is a temporary necessity during the controller replacement process to avoid potential loss of access to the data.

About this task

After resetting the authentication keys to an MSID key, the data on the disks is no longer protected by secret authentication keys. You must verify the physical safety of the disks during the replacement or upgrade process. The steps for Storage Encryption are also required for Volume Encryption.

Steps

1. Access the nodeshell:

   system node run -node *node_name*

2. Display the key ID for each self-encrypting disk on the original system:

   disk encrypt show

   Example

   ```
   disk encrypt show
   Disk       Key ID                                                           Locked?
   0c.00.1    0x0                                                              No
   0c.00.0    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.3    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.4    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.2    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.5    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   ```

   The first disk in the example is associated with an MSID key; the other disks are associated with a non-MSID key.

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

   - Rekey the disks individually, once for each disk:

     disk encrypt rekey 0x0 *disk_name*

   - Rekey all of the disks at once:

     disk encrypt rekey 0x0 *

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

   disk encrypt show

   Example

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

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

5. Exit the nodeshell and return to the clustershell:

   exit

6. Repeat step *1* on page 34 through step *5* on page 35 for each individual node or HA pair.

Shutting down the target controller

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

Choices

- Shutting down a node running ONTAP on page 35

Shutting down a node running ONTAP

To shut down an impaired node, you must determine the status of the node and, if necessary, take over the node so that the healthy node continues to serve data from the impaired node storage.

About this task

You must leave the power supplies turned on at the end of this procedure to provide power to the healthy node.

Steps

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

   a. Change the privilege level to advanced, entering **y** when prompted to continue: **set -privilege advanced**

      The advanced prompt (*>) appears.

   b. Verify the status of the node members in the cluster: **cluster show -epsilon ***

      #### Example

      The following example displays information about the health and eligibility of the nodes in the cluster:

      ```
      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: You must not assign epsilon to a node that has to be replaced.

      Note: In a cluster with a single HA pair, you must not assign epsilon to either node.

   c. Perform 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	**a.** Exit advanced mode: **set -privilege admin** **b.** Proceed to Step 3.

If...	Then...
All nodes show **`true`** for both health and eligibility, and epsilon is assigned to the impaired node	**a.** Remove epsilon from the node: **`cluster modify -node node1 -epsilon false`** **b.** Assign epsilon to a node in the cluster: **`cluster modify -node node4 -epsilon true`** **c.** Exit advanced mode: **`set -privilege admin`** **d.** Go to Step 3.
The impaired node shows **`false`** for health and is the epsilon node	**a.** Change the privilege level to advanced: **`set -privilege advanced`** **b.** Remove epsilon from the node: **`cluster modify -node node1 -epsilon false`** **c.** Assign epsilon to a node in the cluster: **`cluster modify -node node4 -epsilon true`** **d.** Exit advanced mode: **`set -privilege admin`** **e.** Proceed to the next step.
The impaired node shows **`false`** for health and is not the epsilon node	**a.** Proceed to the next step.
Any node shows **`false`** for eligibility	**a.** Resolve any cluster issues as needed. **b.** Exit advanced mode: **`set -privilege admin`**
Any node other than the impaired node shows **`false`** for health	**a.** Correct the problems that caused the health issues on the nodes. **b.** Exit advanced mode: **`set -privilege admin`**

2. If the impaired node is part of an HA pair, disable the `auto-giveback` option from the console of the healthy node:
`storage failover modify -node local -auto-giveback false`

3. Bring the impaired node to the LOADER prompt:

If the impaired node is in...	Then...
A stand-alone configuration and is running	Halt the impaired node: **`system -node halt impaired_node_name`**
A stand-alone configuration and is not running and is not at the LOADER prompt	Resolve any issues that caused the node to quit running, power-cycle it, and then halt the boot process by entering **`Ctrl-C`** and responding **`y`** to take the node to the LOADER prompt.

If the impaired node is in...	Then...
An HA pair	The impaired node is at the LOADER prompt, it is ready for service. Otherwise, take one of the following actions, as applicable: • If the impaired node shows the ONTAP prompt, then take over the impaired node from the healthy node and be prepared to interrupt the reboot: **`storage failover takeover -ofnode impaired_node_name`** When prompted to interrupt the reboot, you must press **`Ctrl-C`** to go to the LOADER prompt. • If the display of the impaired node shows the `Waiting for giveback` message, then press **`Ctrl-C`** and respond **`y`** to take the node to the LOADER prompt. • If the impaired node does not show either the `Waiting for giveback` message or an ONTAP prompt, then power-cycle the node. You must contact technical support if the node does not respond to the power cycle.

4. Respond to the applicable wizard:

If LED is...	Then...
Off	NVRAM has no data. **Note:** You can power down and disconnect the battery.
Flashing	NVRAM is destaging. **Note:** You must wait for two minutes to complete the destaging operation.
On	NVRAM has data. **Note:** If you are sure that the data stored in NVRAM is not required, then you can proceed with shut down and replacement. If you have to wipe the NVRAM memory on the impaired controller module, then you must contact technical support for instructions.

5. Shut down the impaired node.

   Note: If the node is in an HA pair, the impaired node console should show the `Waiting for giveback...` message.

   The method that you use to shut down the node depends on whether you use remote management through the node's Service Processor (SP):

If the SP is...	Then...
Configured	Log in to the SP of the impaired node, and then turn off the power: **`system power off`** .
Not configured	At the impaired node prompt, press **`Ctrl-C`** and respond **`y`** to halt the node.

6. If the system is in a stand-alone configuration, shut down the power supplies, and then unplug both of the power cords from the power source.

Verifying the new controller module has no content in NVRAM

You must check that the new controller module has no content in NVRAM before completing the replacement.

Steps

1. Check the NVRAM LED on the controller module.

You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the nonvolatile memory (NVMEM). The NVMEM LED is marked with a battery symbol and is located about 3" to the left of the label showing the MAC address on the controller module.

2. If the NVRAM LED is not flashing, there is no content in the NVRAM; You can skip the following steps and proceed to the next task in this procedure.
3. If the NVRAM LED is flashing, there is data in the NVRAM and you must disconnect the battery to clear the memory:
   a. If you are not already grounded, properly ground yourself.
   b. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and unplug the battery cable from the socket.

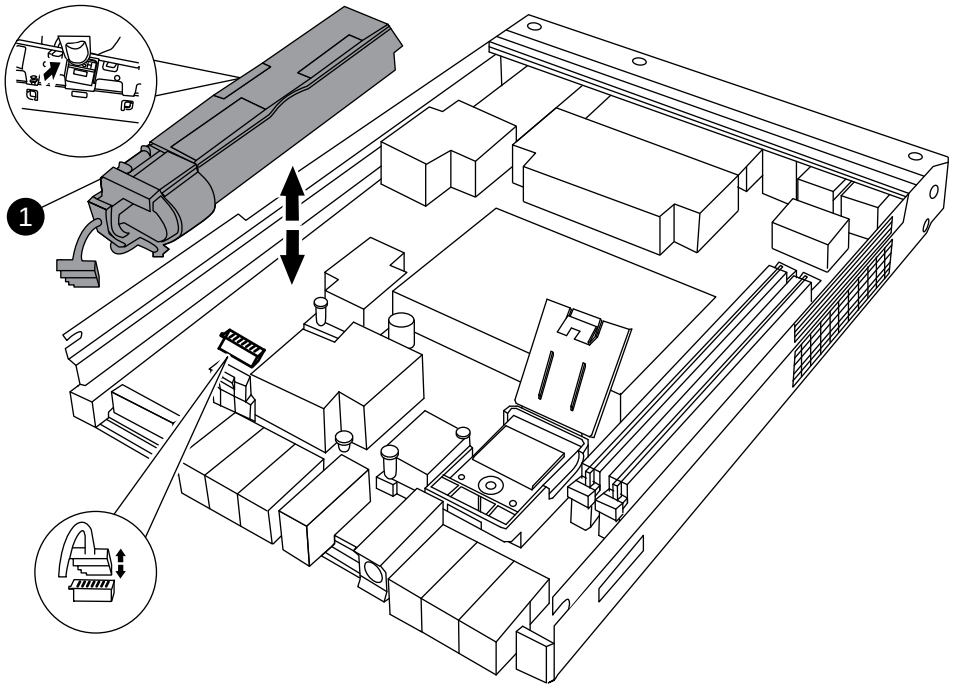

1	NVMEM battery

   c. Confirm that the NVRAM LED is no longer lit.
   d. Reconnect the battery connector.
4. Return to step 1 of this procedure to recheck the NVRAM LED.

Replacing the controller module hardware

To replace the controller module hardware, you must remove the impaired node, move FRU components to the replacement controller module, install the replacement controller module in the chassis, and then boot the system to Maintenance mode.

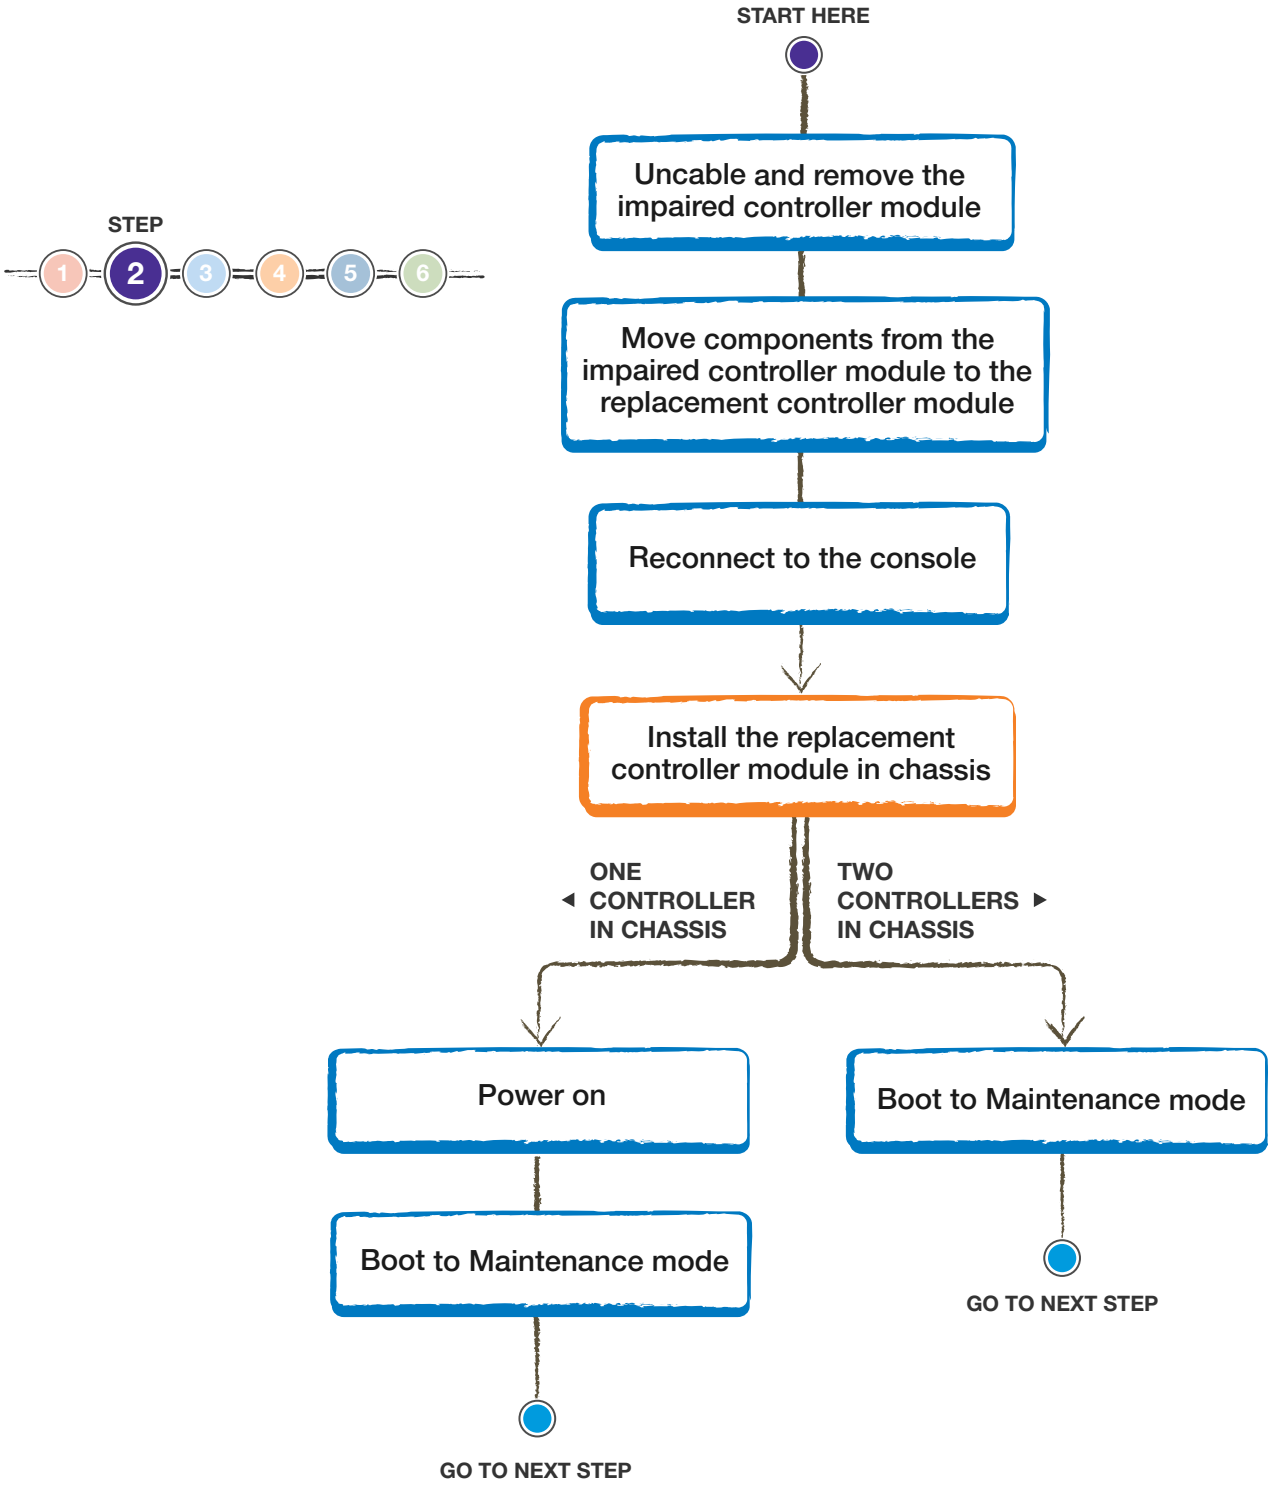

Steps

1. Removing the controller module and moving the components on page 40
2. Installing the new controller module and booting the system on page 45

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 for two minutes after takeover of the impaired node to confirm 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.

Steps

1. Removing the controller module from the system on page 40
2. Moving the boot device on page 42
3. Moving the NVRAM battery on page 43
4. Moving the DIMMs to the new controller module on page 43

Removing the controller module from the system

To replace the controller module, you must first remove the old controller module from 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 arm, 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 arm so that when you reinstall the cable management arm, the cables are organized.
3. Remove the cable management arms from the left and right sides of the controller module and set them aside.

   The illustration shows the cable management arms on a FAS2552 system. The procedure is the same for all FAS25xx systems.

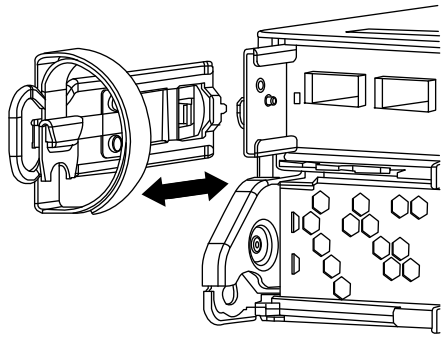

4. Squeeze the latch on the cam handle until it releases, as shown in the following illustration. Open the cam handle fully to release the controller module from the midplane, and then, using two hands, pull the controller module out of the chassis.

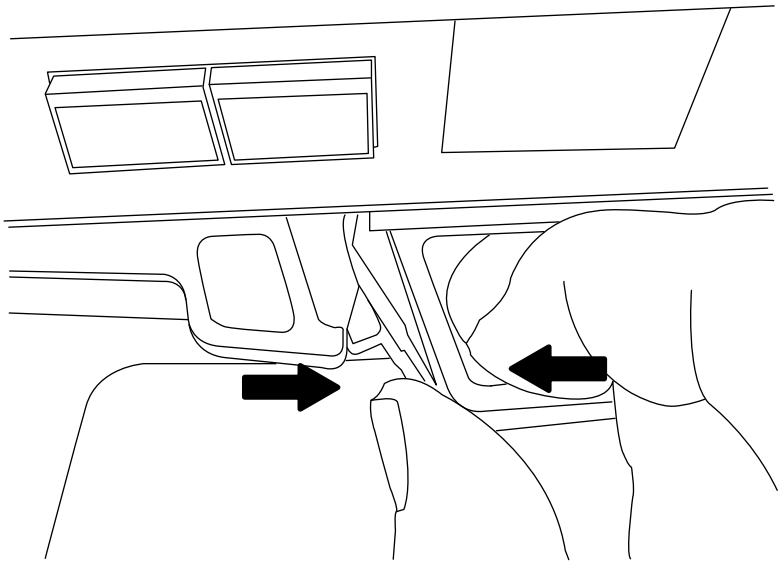

5. Turn the controller module over and open it by pressing the button to release the cover, and then slide the cover out.

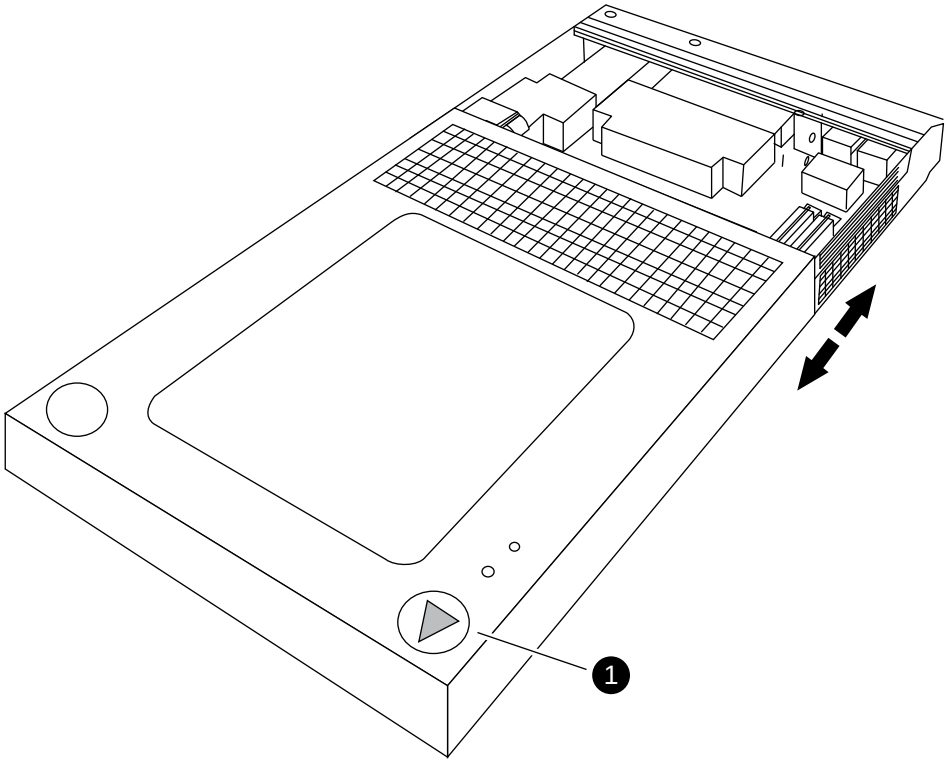

1	Button to release controller module cover

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:

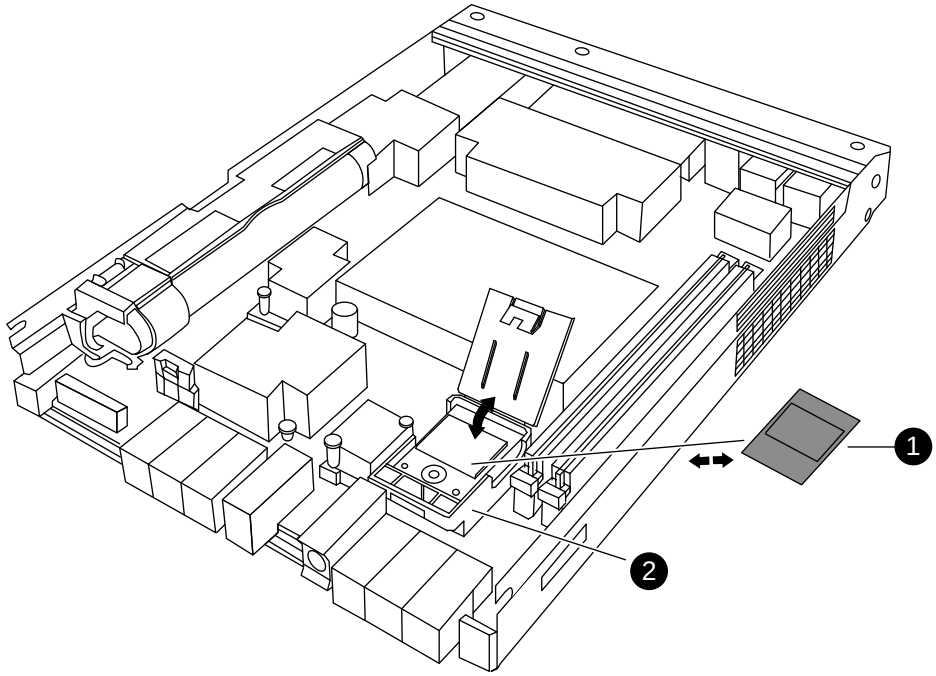

1	Boot device
2	Boot device holder; not removable

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, 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. Verify that the boot device is seated squarely and completely in the socket or connector.

   If necessary, remove the boot device and reseat it into the socket.

6. Close the boot device cover.

Moving the NVRAM battery

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

Steps

1. Locate the battery, press the clip on the face of the battery plug to release the lock clip from the plug socket, and then unplug the battery cable from the socket.

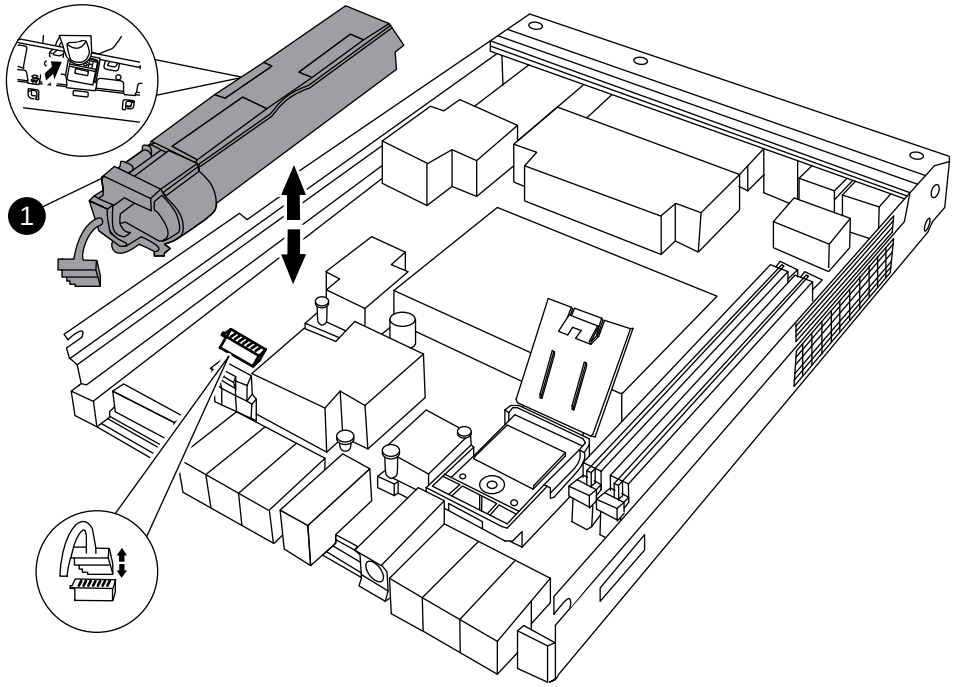

1	NVMEM battery

2. Grasp the battery and press the tab marked PUSH, and then lift the battery out of the holder and controller module.
3. In the new controller module, seat the battery in the holder.

   Attention: Do not connect the NVRAM keyed battery plug into the socket until after the NVRAM DIMM has been installed.

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. Verify that the NVMEM battery cable connector is not plugged into the socket.
2. Locate the DIMMs.

   If you are moving DIMMs on a FAS25xx system:

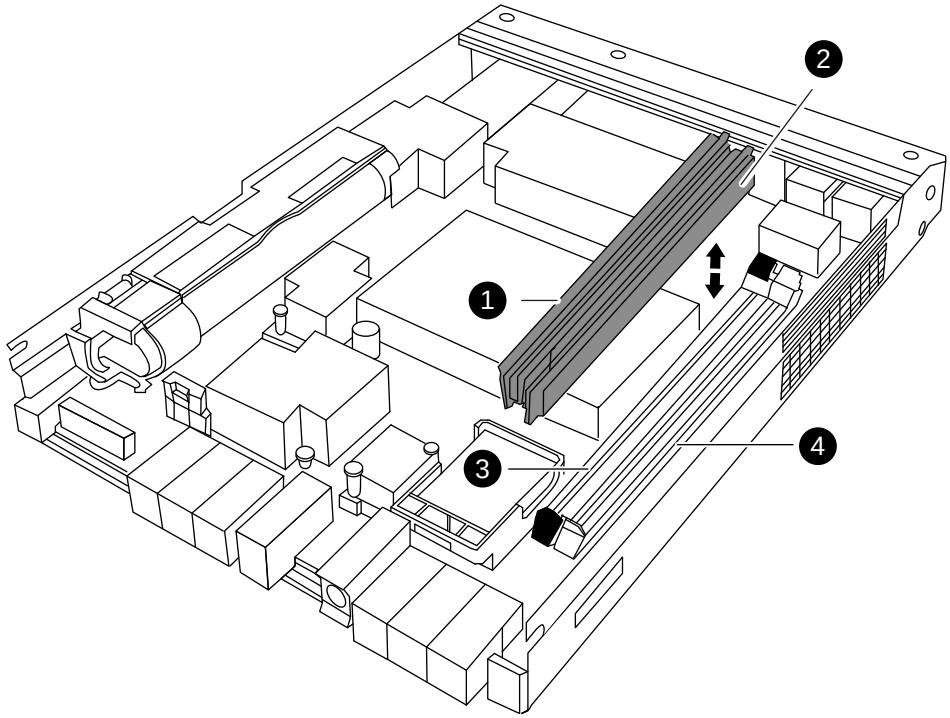

1	System DIMM
2	NVMEM DIMM The NVMEM DIMM has an *NVMEM* label on one of the chips.
3	System DIMM slot
4	NVMEM DIMM slot The NVMEM DIMM slot has white ejector tabs.

3. 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.
4. 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.

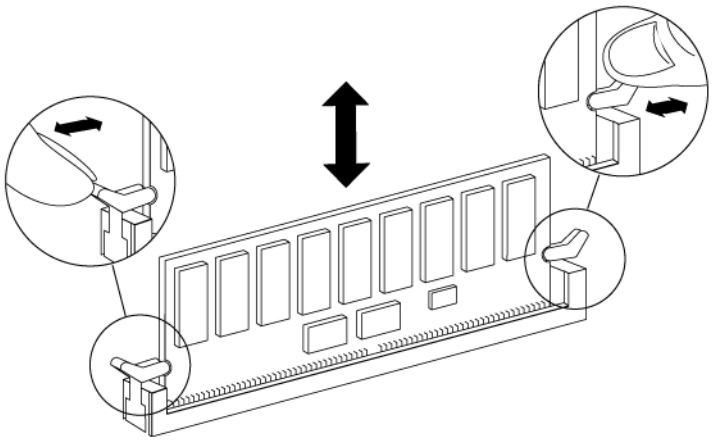

Attention: You must carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

5. Locate the corresponding slot for the DIMM in the new controller module, align the DIMM over the slot, and then 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, you should realign the DIMM with the slot and reinsert it.

   Important: You must install the NVMEM DIMM only in the NVMEM DIMM slot.

6. 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.

7. 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.

8. Repeat these steps to move additional DIMMs, as required.

9. In the new controller module, orient the NVMEM battery cable connector to the socket on the controller module and plug the cable into the socket.

   You must ensure that the plug locks down onto the socket on the controller module.

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 might update the system firmware when it boots. You must 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: You must not completely insert the controller module in the chassis until instructed to do so.

2. Recable the management port or serial console port so that 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	**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:** You must not use excessive force when sliding the controller module into the chassis; you might damage the connectors. **b.** Enter one of the following commands from the healthy node's console and wait for the giveback to complete: For systems operating in... / Issue the command... 7-Mode / `cf giveback` Clustered Data ONTAP / • In Data ONTAP 8.2 or later: `storage failover giveback –ofnode impaired_node_name` **c.** If you have not already done so, reinstall the cable management arm, and then tighten the thumbscrew on the cam handle on back of the controller module. **d.** Bind the cables to the cable management device with the hook and loop strap.
A stand-alone configuration	**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:** You must 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. **c.** If you have not already done so, reinstall the cable management arm, and then tighten the thumbscrew on the cam handle on back of the controller module. **d.** 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 an HA configuration you must confirm that the healthy node remains down.

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

Restoring and verifying the system configuration after hardware replacement

After replacing the hardware components, you should verify the low-level system configuration of the replacement controller and reconfigure FC settings if necessary.

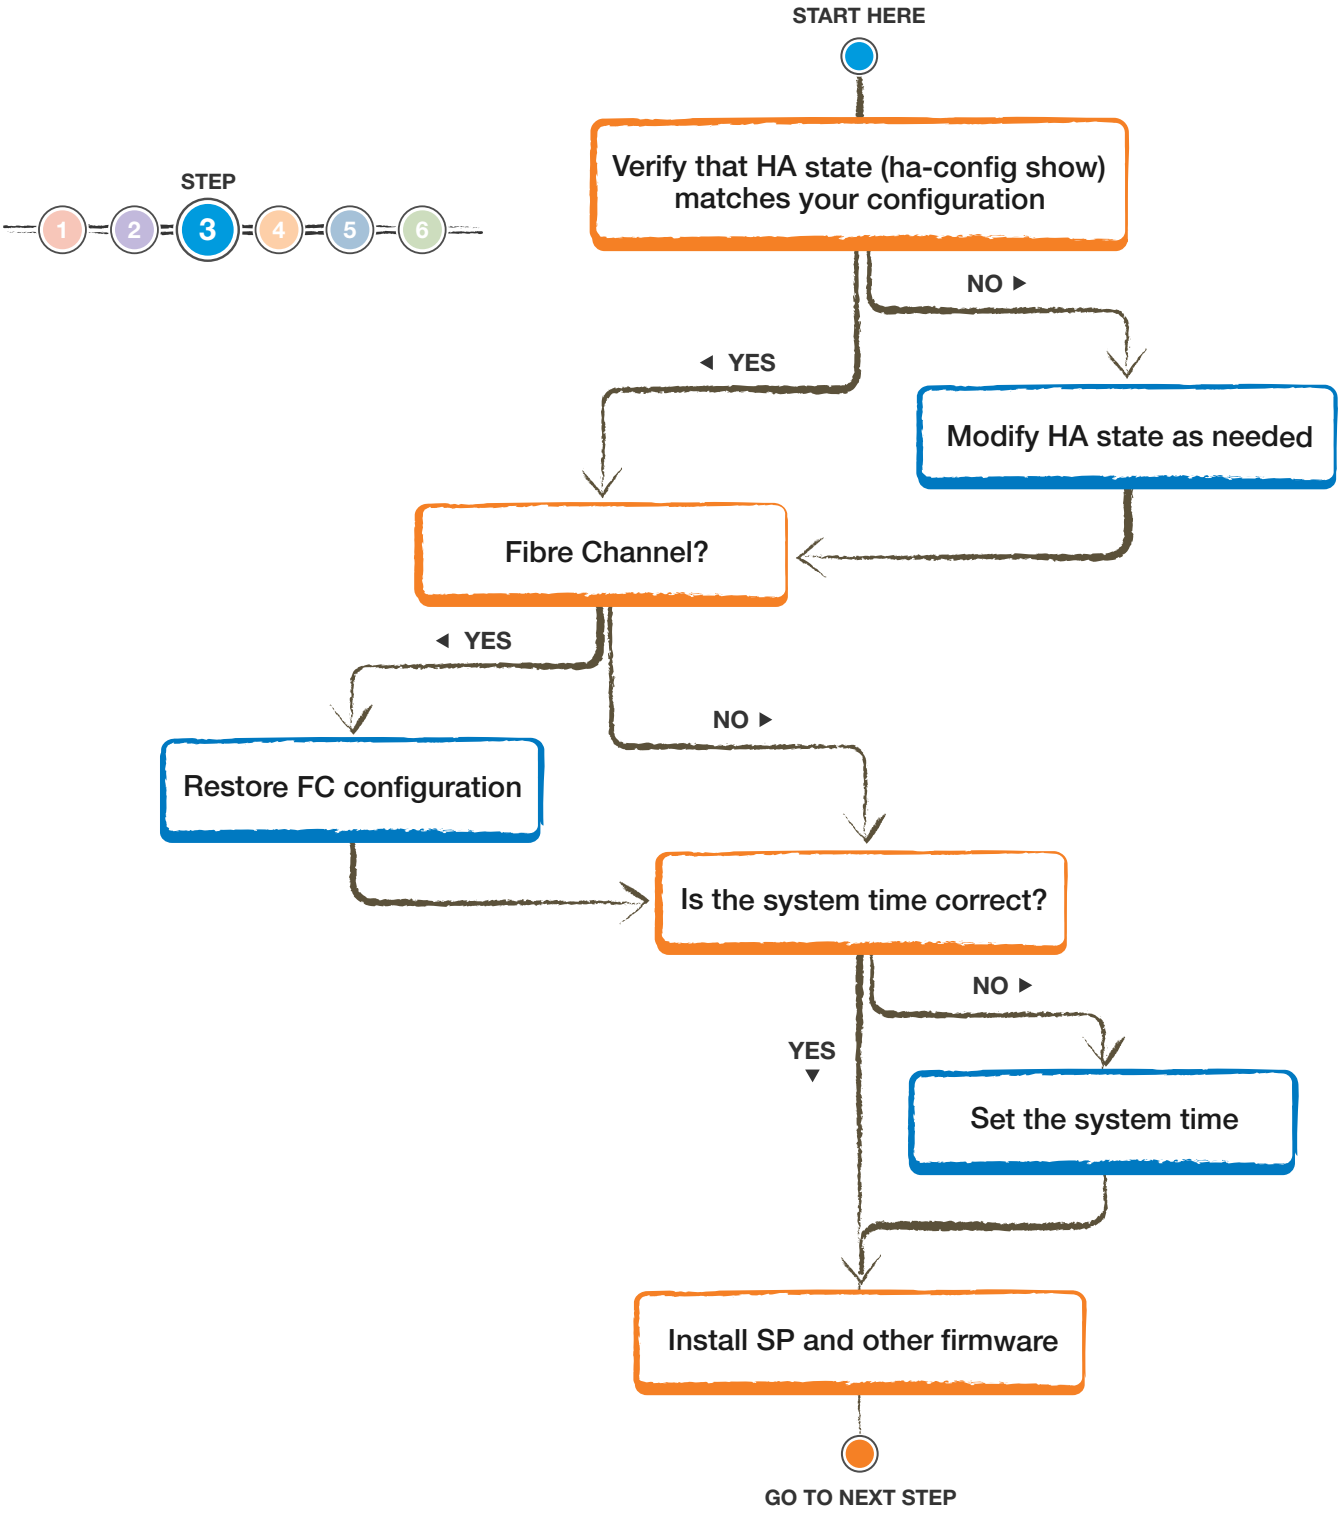

Steps

1. Verifying and setting the HA state of the controller module on page 48
2. Restoring Fibre Channel configurations (CNA) on page 48
3. Restoring 10 Gb Ethernet configurations (CNA) on page 49
4. Setting the system time after replacing the controller module on page 50
5. Installing the firmware after replacing the controller module on page 50

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, 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 Fibre Channel configurations (CNA)

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

Before you begin

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

Steps

1. Verify the FC port configuration:

ucadmin show

2. In Maintenance mode, restore the FC port configuration:

   ucadmin modify -mode fc -type *initiator|target adapter_name*

   - *initiator* is entered if you are connecting to a Fibre Channel tape device.
   - *target* is entered if you are in a SAN configuration.

3. Take one of the following actions, depending on your configuration:

If the FC port configuration is...	Then...
The same for both ports	Answer **y** when prompted by the system because modifying one port in a port pair modifies the other port as well.
Different	a. Answer **n** when prompted by the system. b. Restore the FC port configuration in 7-Mode only: **ucadmin modify -mode fc -type *initiator|target adapter_name*** c. Restore the FC port configuration in CDOT only: **unified-connect modify-mode fc -type** *initiator|targetadapter_name*

4. Exit Maintenance mode:

   halt

   After you issue the command, wait until the system stops at the LOADER prompt.

5. Boot the node back into Maintenance mode for the configuration changes to take effect.

6. Verify the values of the variables:

   ucadmin show

Restoring 10 Gb Ethernet configurations (CNA)

Because the onboard Converged Network Adapter (CNA) ports are not preconfigured as 10 Gb Ethernet, you must restore any 10 Gb Ethernet port configurations in your HA pair before you bring the node back into service; otherwise, you might experience a disruption in service.

Before you begin

You must have the values of the 10 Gb Ethernet port settings that you saved earlier.

Steps

1. In Maintenance mode, program the Ethernet ports in 7-mode only:

   ucadmin modify -mode cna *adapter_name*

2. Because modifying one port in a port pair modifies the other port, answer **y** when prompted by the system.

3. Exit Maintenance mode:

   halt

   After you issue the command, wait until the system stops at the LOADER prompt.

4. Boot the node back into Maintenance mode for the configuration changes to take effect.

5. Verify the values of the variables:

ucadmin show

Setting the system time after replacing the controller module

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, verify 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. Determine the system time by using the `date` command on the healthy node (if the system is in an HA pair) or another reliable time source.
3. Set the date in GMT on the replacement node:

   set date *mm/dd/yyyy*
4. Set the time in GMT on the replacement node:

   set time *hh*:*mm*:*ss*

Installing the firmware after replacing the controller module

After replacing the controller module, you must install the latest firmware if your system is running a version of Data ONTAP earlier than 8.2, 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

If your system is running ONTAP 8.2 or later, the SP firmware and BIOS automatically update to the baseline image included with the ONTAP version. Other system firmware from the old controller module still resides on the boot device and typically does not need updating.

If your system is running ONTAP 8.2 or later, you should skip this procedure.

Steps

1. Check the configuration of the SP from the LOADER prompt:

   sp status

   For the latest release of SP firmware, log in to the NetApp Support Site at *mysupport.netapp.com* and update it, if needed, in the following steps.
2. Log in to the SP from an administration host:

   ssh *username@SP_IP_address*
3. Download and install the most current version of firmware for your system by 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.

Related information

Find a System Administration Guide for your version of ONTAP 9
Find a System Administration Guide for your version of Data ONTAP 8

Running diagnostics tests after replacing a controller module

You should run focused diagnostic tests for specific components and subsystems whenever you replace a component of the controller.

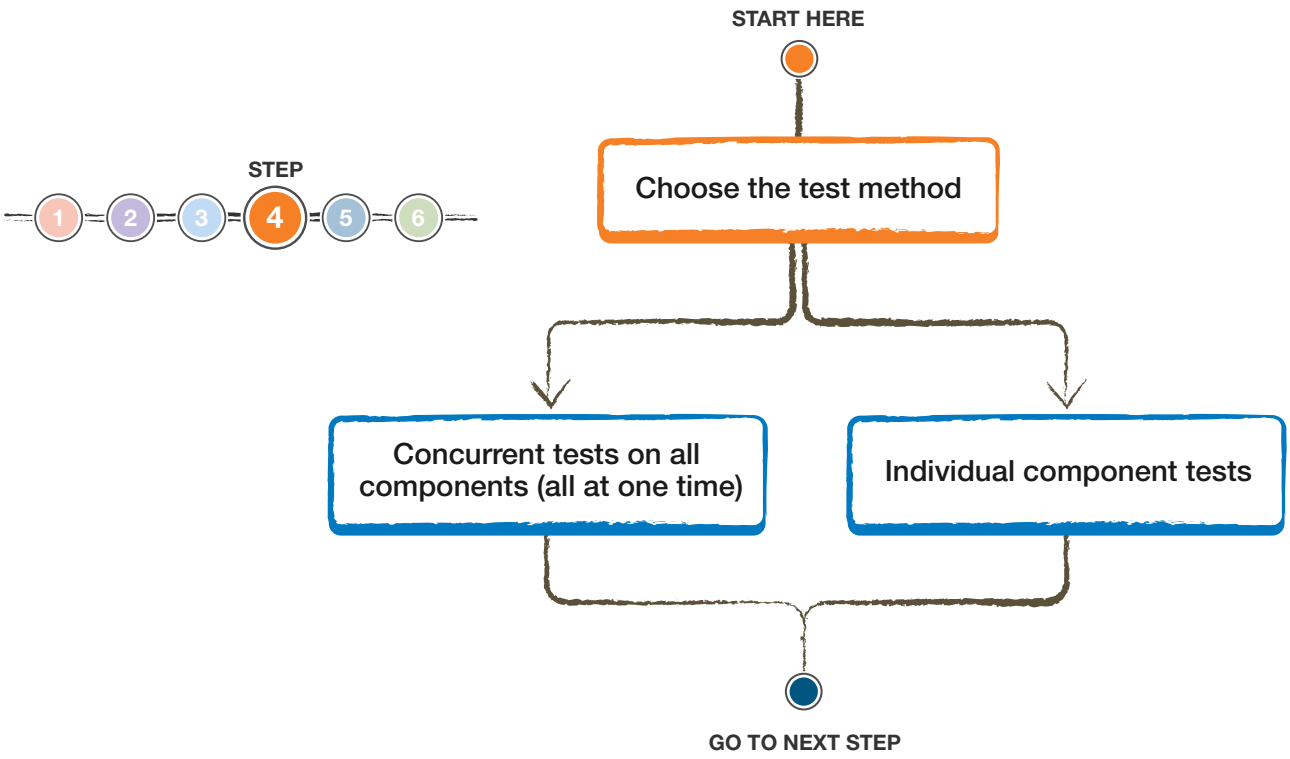

Before you begin

- Your system must be at the LOADER prompt to start system-level diagnostics.
- For ONTAP 8.2 and later, you do not require 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, run the system-level diagnostic test: **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 that are designed specifically for system-level diagnostics.

   Important: During the boot_diags process, you might see a prompt warning that when entering Maintenance mode in an HA configuration, you must confirm that the partner remains down. To continue to Maintenance mode, you should enter **y**

3. Clear the status logs: **sldiag device clearstatus**

4. Display and note the available devices on the controller module: **sldiag device show -dev mb**

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

   - **bootmedia** is the system booting device.
   - **cna** is a Converged Network Adapter or interface that is not connected to a network or storage device.
   - **env** is the motherboard environmentals.
   - **mem** is the system memory.
   - **nic** is a network interface card.
   - **nvmem** is a hybrid of NVRAM and system memory.
   - **sas** is a Serial Attached SCSI device that is not connected to a disk shelf.

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

Choices

- Running diagnostics tests concurrently after replacing the controller module on page 52
- Running diagnostics tests individually after replacing the controller module on page 53

Running diagnostics tests concurrently after replacing the controller module

After replacing the controller module, you can 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 that you make. If you run more tests in addition to the default tests, the diagnostic test process takes longer to complete.

Steps

1. Display and note the available devices on the controller module: **sldiag device show -dev mb**

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

   - **bootmedia** is the system booting device.
   - **cna** is a Converged Network Adapter or interface that is not connected to a network or storage device.
   - **env** is the motherboard environmentals.
   - **mem** is the system memory.
   - **nic** is a network interface card.
   - **nvmem** is a hybrid of NVRAM and system memory.

- **sas** is a Serial Attached SCSI device that is not connected to a disk shelf.

2. Review the enabled and disabled devices in the output from step 1 and then determine which tests you want to run concurrently.
3. List the individual tests for each device:

   sldiag device show -dev *dev_name*
4. Verify that the tests were modified: **sldiag device show**
5. Repeat steps 2 through 4 of this procedure for each device.
6. Run diagnostics on all the devices: **sldiag device run**

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

   The tests are complete when the following message is displayed:

   ```
   *> <SLDIAG:_ALL_TESTS_COMPLETED>
   ```
7. After the tests are complete, verify that there are no hardware problems on your storage system:

   sldiag device status -long -state failed
8. Correct any issues that are found, and repeat this procedure.

Running diagnostics tests individually after replacing the controller module

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

Steps

1. Clear the status logs: **sldiag device clearstatus**
2. Display the available tests for the selected devices:

Device type	Command
boot media	**sldiag device show -dev bootmedia**
cnafcal	**sldiag device show -dev cna** **sldiag device show -dev fcal**
env	**sldiag device show -dev env**
mem	**sldiag device show -dev mem**
nic	**sldiag device show -dev nic**
nvmem	**sldiag device show -dev nvmem**
sas	**sldiag device show -dev sas**

3. 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.

4. 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

5. Run the selected tests:

Device type	Command
boot media	sldiag device run -dev bootmedia
cnafcal	sldiag device run -dev cna sldiag device run -dev fcal
env	sldiag device run -dev env
mem	sldiag device run -dev mem
nic	sldiag device run -dev nic
nvmem	sldiag device run -dev nvmem
sas	sldiag device run -dev sas

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

```
<SLDIAG:_ALL_TESTS_COMPLETED>
```

6. Verify that no tests failed:

Device type	Command
boot media	sldiag device status -dev bootmedia -long -state failed
cnafcal	sldiag device status -dev cna sldiag device status -dev fcal
env	sldiag device status -dev env -long -state failed
mem	sldiag device status -dev mem -long -state failed
nic	sldiag device status -dev nic -long -state failed
nvmem	sldiag device status -dev nvmem
sas	sldiag device status -dev sas -long -state failed

   Any tests that failed are displayed.

7. 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: **`sldiag device clearstatus`** **b.** Verify that the log is cleared: **`sldiag device status`** The following `SLDIAG: No log messages are present.` default response is displayed. You have completed system-level diagnostics.
Resulted in some test failures	Determine the cause of the problem: **a.** Exit Maintenance mode: **`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 then unplug them from the power sources. **c.** Check the controller module you are servicing to verify that you have observed all of 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 then turn them on. **e.** Select **Boot to Maintenance mode** from the menu. **f.** Exit Maintenance mode: **`halt`** After you issue the command, you must wait until the system stops at the LOADER prompt. **g.** Enter **`boot_diags`** at the prompt, and then rerun the system-level diagnostic test.

8. Exit system-level diagnostics, and continue with recabling and restoration of the storage system.

Completing the recabling and final restoration of operations

To complete the replacement procedure, you must recable the storage system, confirm disk reassignment, restore the NetApp Storage Encryption configuration (if necessary), and install licenses for the new controller.

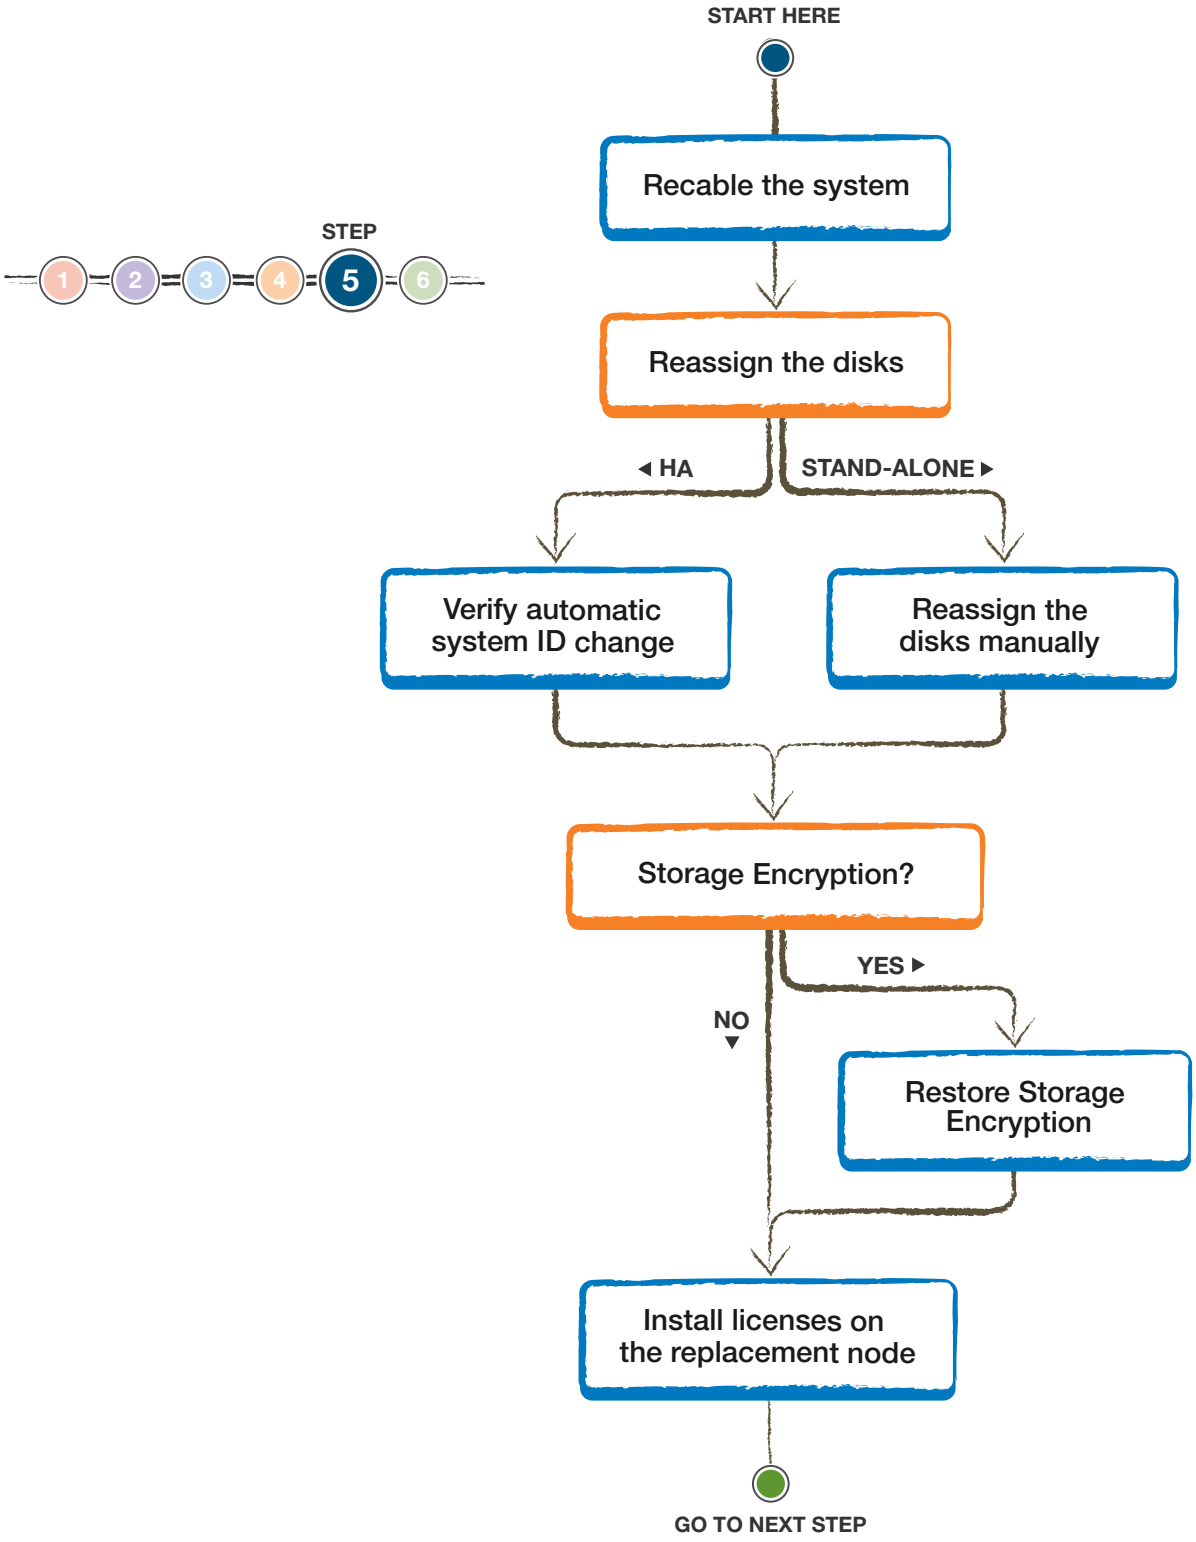
STEP
1
2
3
4
5
6
START HERE
Recable the system
Reassign the disks
◀ HA
STAND-ALONE ▶
Verify automatic system ID change
Reassign the disks manually
Storage Encryption?
YES ▶
NO ▼
Restore Storage Encryption
Install licenses on the replacement node
GO TO NEXT STEP

Steps

1. Recabling the system on page 57
2. Reassigning disks on page 57
3. Installing licenses for the replacement node in clustered Data ONTAP on page 62
4. Restoring Storage and Volume Encryption functionality on page 63
5. Verifying LIFs and registering the serial number on page 63

Recabling the system

After running diagnostics, you must recable the storage and network connections of the controller module.

Steps

1. Reinstall the cable management arms and recable the controller module, 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:

      ToolChest

      The “Quick Start Guide” provides instruction to collect and analyze data from your system.

      Config Advisor Quick Start Guide

   b. Check the rules for “SAS Cabling Checks” and then examine the output from Config Advisor.

      You must verify that all of the disk shelves are displayed and that all disks appear in the output. You must correct any cabling issues that you might find.

Related information

Disk Shelves Documentation

Reassigning disks

If the storage system is in an HA pair, the system ID of the new controller module is automatically assigned to the disks when the giveback occurs at the end of the procedure. In a stand-alone system, you must manually reassign the ID to the disks.

About this task

You must use the correct procedure for your configuration:

Controller redundancy	Then use this procedure...
HA pair	*Verifying the system ID change on a system operating in clustered Data ONTAP* on page 57
Stand-alone	*Manually reassigning the system ID on a stand-alone system in clustered Data ONTAP* on page 61

Verifying the system ID change on an HA system running clustered Data ONTAP

If you are running ONTAP 8.2 or later, you must confirm the system ID change when you boot the replacement node, and then verify that the change was implemented.

About this task

This procedure applies only to systems running clustered Data ONTAP in an HA pair.

Steps

1. If the replacement node is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode:

   `halt`

   After you issue the command, you must wait until the system stops at the LOADER prompt.

2. If you are running Data ONTAP 8.2.2 or earlier at the replacement node prompt, set the environmental variables:

   a. Confirm that the new controller module boots in clustered Data ONTAP: **`setenv bootarg.init.boot_clustered true`**

   b. Save the settings: **`saveenv`**

3. From the LOADER prompt on the replacement node, boot the node:

If you are running ONTAP...	Then...
8.2.x and earlier	**a.** Boot the node: **`boot_ontap`** **b.** Press **`Ctrl-c`** when prompted to display the boot menu.
8.3 and later	Boot the node: **`boot_ontap`**

   If you are prompted to override the system ID due to a system ID mismatch, enter **`y`**.

4. Wait until the `Waiting for giveback...` message is displayed on the replacement node console and then, on the healthy node, verify that the controller module replacement has been detected and the new partner system ID has been automatically assigned.

   ### Example

```
node1::*> storage failover show
                                 Takeover
Node              Partner        Possible     State Description
------------      ------------   --------     -----------------------------------
node1             node2          false        System ID changed on partner (Old:
                                               151759755, New: 151759706), In
takeover
node2             node1          -            Waiting for giveback (HA mailboxes)
```

5. From the healthy node, verify that any coredumps are saved:

   a. Change the privilege level to advanced, entering **`y`** when prompted to continue:

      `set -privilege advanced`

      The advanced prompt (`*>`) appears.

   b. Save any coredumps:

      `system node run -node local-node-name partner savecore`

   c. Wait for the savecore to finish before issuing the giveback.

      If desired, monitor the progress of the savecore command:

      `system node run -node local-node-name partner savecore -s`

   d. Change the privilege level back to admin:

```
set -privilege admin
```

6. Your next step depends on the version of ONTAP your system is running.

If your system is running...	Then...
Data ONTAP 8.2.0 and earlier or ONTAP 8.2.2 and later	Go to the next step.
Data ONTAP 8.2.1	Disable automatic takeover on reboot from the healthy node: **`storage failover modify -node replacement-node-name -onreboot false`**

7. Your next step depends on your version of ONTAP:

If your system is running...	Then...
Data ONTAP 8.2.0 and earlier or ONTAP 8.2.2 and later	Complete the following substeps after the replacement node is displaying the `Waiting for Giveback...` message: **a.** Give back the node: **`storage failover giveback -ofnode replacement_node_name`** As the replacement node boots up, it might again display the prompt warning of a system ID mismatch and asking to override the system ID. You can respond **`Y`**. The replacement node takes back its storage and finishes booting up, and then reboots and is again taken over by the healthy node. As the replacement node boots up the second time, it might again display the prompt warning of a system ID mismatch and asking to override the system ID. You can respond **`Y`**. **b.** After the node displays `Waiting for Giveback...`, give back the node: **`storage failover giveback -ofnode replacement_node_name`** As the replacement node boots up, it might again display the prompt warning of a system ID mismatch and asking to override the system ID. You can respond **`Y`**. The replacement node takes back its storage and finishes booting up to the ONTAP prompt. **Note:** If the giveback is vetoed, you can consider overriding the vetoes. *ONTAP 9 High-Availability Configuration Guide* **c.** Monitor the progress of the giveback operation: `storage failover show-giveback` **d.** Wait until the `storage failover show-giveback` command output indicates that the giveback operation is complete. **e.** Confirm that the HA pair is healthy and that takeover is possible by using the `storage failover show` command. The output from the `storage failover show` command should not include the `System ID changed on partner` message.

If your system is running...	Then...
Data ONTAP 8.2.1 only	Complete the following substeps after the replacement node is displaying the `Waiting for Giveback...` message: **a.** Give back the node: **`storage failover giveback -ofnode replacement_node_name`** As the replacement node boots up, it might display the prompt warning of a system ID mismatch and asking to override the system ID. You can respond **`Y`**. The replacement node takes back its storage, finishes booting up, and then reboots. **b.** Manually take over the replacement node: **`storage failover takeover -ofnode replacement_node_name`** As the replacement node boots up the second time, it might again display the prompt warning of a system ID mismatch and asking to override the system ID. You can respond **`Y`**. **c.** Give back the node: **`storage failover giveback -ofnode replacement_node_name`** As the replacement node boots up, it might again display the prompt warning of a system ID mismatch and asking to override the system ID. You can respond **`Y`**. The replacement node takes back its storage and finishes booting up to the Data ONTAP prompt. **Note:** If the giveback is vetoed, you can consider overriding the vetoes. *ONTAP 9 High-Availability Configuration Guide* **d.** Monitor the progress of the giveback operation by using the `storage failover show-giveback` command. **e.** Wait until the `storage failover show-giveback` command output indicates that the giveback operation is complete. **f.** Confirm that the HA pair is healthy and that takeover is possible by using the `storage failover show` command. The output from the `storage failover show` command should not include the `System ID changed on partner` message.

8. Verify that the disks or FlexArray Virtualization LUNs were assigned correctly: **`storage disk show -ownership`**

Example

The disks belonging to the replacement node should show the new system ID for the replacement node. In the following example, the disks owned by node1 now show the new system ID, 1873775277:

```
node1> storage disk show -ownership

Disk  Aggregate Home  Owner  DR Home  Home ID    Owner ID  DR Home ID Reserver  Pool
----- ------    ----- ------ -------- -------    -------    -------  ---------  ---
1.0.0  aggr0_1  node1 node1  -        1873775277 1873775277  -       1873775277 Pool0
1.0.1  aggr0_1  node1 node1           1873775277 1873775277  -       1873775277 Pool0
.
.
.
```

9. Verify that the expected volumes are present for each node: **`vol show -node node-name`**

10. If you disabled automatic takeover on reboot, reenable it on the healthy node console: **`storage failover modify -node replacement-node-name -onreboot true`**

Manually reassigning the system ID on a stand-alone system in clustered Data ONTAP

In a stand-alone system, you must manually reassign disks to the new controller's system ID and set the bootarg.init.boot_clustered bootarg before you return the system to normal operating condition.

About this task

This procedure applies only to systems that are running Data ONTAP operating in 7-Mode or are stand-alone.

Steps

1. If you have not already done so, reboot the replacement node, interrupt the boot process by entering **Ctrl-C**, and then select the option to boot to Maintenance mode from the displayed menu.

   You must enter **Y** when prompted to override the system ID due to a system ID mismatch.

2. View the system IDs:

   disk show -a

   Note: Make a note of the old system ID, which is displayed as part of the disk owner column.

   #### Example

   The following example shows the old system ID of 118073209:

```
*> disk show -a
Local System ID: 118065481

 DISK     OWNER        POOL    SERIAL NUMBER   HOME
--------  -----------  ------  -------------   -------------
system-1  (118073209)  Pool0   J8XJE9LC        system-1  (118073209)
system-1  (118073209)  Pool0   J8Y478RC        system-1  (118073209)
.
.
.
```

3. Reassign disk ownership by using the system ID information obtained from the `disk show` command:

   disk reassign -s *old system ID*

   In the case of the preceding example, the command is **disk reassign -s 118073209**.

   You can respond **Y** when prompted to continue.

4. Verify that the disks were assigned correctly: **disk show -a**

   You must verify that the disks belonging to the replacement node show the new system ID for the replacement node. In the following example, the disks owned by system-1 now show the new system ID, 118065481:

   #### Example

```
*> disk show -a
Local System ID: 118065481

  DISK     OWNER          POOL   SERIAL NUMBER  HOME
-------    -------------  -----  -------------  -------------
system-1  (118065481)     Pool0  J8Y0TDZC       system-1  (118065481)
system-1  (118065481)     Pool0  J8Y09DXC       system-1  (118065481)
.
.
.
```

5. If the replacement node is in Maintenance mode (showing the `*>` prompt), exit Maintenance mode:

   `halt`

   After you issue the command, you must wait until the system stops at the LOADER prompt.

6. If you are running Data ONTAP 8.2.2 or earlier at the replacement node prompt, set the environmental variables:

   a. Confirm that the new controller module boots in clustered Data ONTAP: **`setenv bootarg.init.boot_clustered true`**

   b. Save the settings: **`saveenv`**

7. Boot the operating system: **`boot_ontap`**

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 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 ONTAP 8.2 and later.

You have a 90-day grace period 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. 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:

   `license clean-up -unused -simulate`

   b. If the list looks correct, remove the unused licenses:

   `license clean-up -unused`

Related information

Find a System Administration Guide for your version of ONTAP 9
Find a System Administration Guide for your version of Data ONTAP 8
NetApp Knowledgebase Answer 1002749: Data ONTAP 8.2 and 8.3 Licensing Overview and References

Restoring Storage and Volume Encryption functionality

After replacing the controller module or NVRAM module for a storage system that you previously configured to use Storage or Volume Encryption, you must perform additional steps to provide uninterrupted Encryption functionality. You can skip this task on storage systems that do not have Storage or Volume Encryption enabled.

Step

1. Restore Storage or Volume Encryption functionality by using the appropriate procedure in the *NetApp Encryption Power Guide*.

   ONTAP 9 NetApp Encryption Power Guide

   Use one of the following procedures, depending on whether you are using onboard or external key management:

   - "Restoring onboard key management encryption keys"
   - "Restoring external key management encryption keys"

Verifying LIFs and registering the serial number

Before returning the replacement node to service, you should verify that the LIFs are on their home ports, and register the serial number of the replacement node if AutoSupport is enabled, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports:

   `network interface show -is-home false`

   If any LIFs are listed as `false`, revert them to their home ports:

   `network interface revert *`

2. Register the system serial number with NetApp Support.

If...	Then...
AutoSupport is enabled	Send an AutoSupport message to register the serial number.
AutoSupport is not enabled	Call NetApp Support to register the serial number.

3. If automatic giveback was disabled, reenable it:

   `storage failover modify -node local -auto-giveback true`

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.

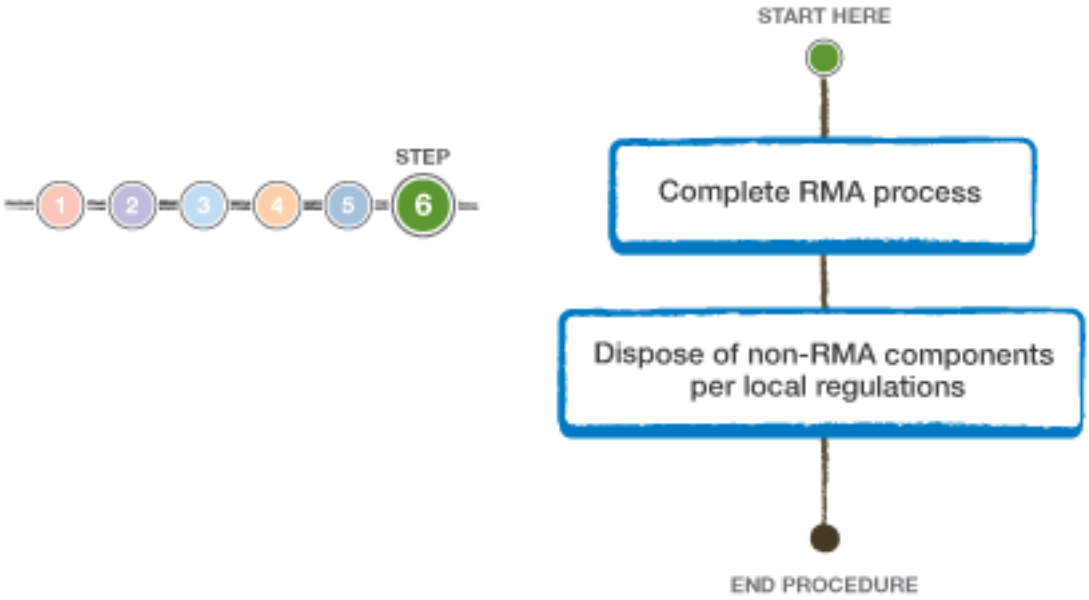

Related information

NetApp Support

Disposing of batteries

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

Related information

https://library.netapp.com/ecm/ecm_download_file/ECMP12475945

Copyright

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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

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.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

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

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277