

Replacing a controller module

To replace a controller module, you must remove the faulty controller module from the system and move the replaceable components from the old controller module to the new controller module. You then install the new controller module in the system chassis.

This procedure applies only to systems running
Data ONTAP[®] 8.1.x
or earlier versions

Before you begin

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

About this task

- This procedure is for systems running Data ONTAP 8.1.x, 8.0.x, or 7.3.x only.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module of the same model type; you cannot upgrade your system by just replacing the controller module.
- You cannot change any disks or disk shelves as part of this procedure.
- The term *system* refers to FAS, V-Series, and SA (FlexCache) systems within this platform family. The procedures apply to all platforms, unless otherwise indicated, except that clustered Data ONTAP procedures do not apply to SA systems.
- It is important that you apply the commands in these steps on the correct systems:
 - The *impaired* node is the node that is being replaced.
 - The *replacement* node is the new node that is replacing the impaired node.
 - The *healthy* node is the surviving node.

Steps

1. [Pre-replacement tasks for systems in SAN configurations](#) on page 2
2. [Shutting down the target controller](#) on page 2
3. [Removing the controller module and moving the components](#) on page 6
4. [Installing the new controller module and booting the system](#) on page 12
5. [Verifying and setting the HA state of the controller module](#) on page 13
6. [Restoring the FC configuration for HA pairs](#) on page 14
7. [Verifying the system time after replacing the controller module in an HA pair](#) on page 15
8. [Installing the firmware after replacing the controller module](#) on page 16

9. [Running diagnostics tests \(controller replacement\)](#) on page 16
10. [Recabling the system](#) on page 20
11. [Reassigning disks](#) on page 21
12. [Performing a final takeover and giveback from the impaired node](#) on page 26
13. [Restoring licenses on the replacement node](#) on page 26
14. [Restoring Storage Encryption functionality after controller module replacement](#) on page 28
15. [Completing the replacement process](#) on page 28

Pre-replacement tasks for systems in SAN configurations

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

Steps

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

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

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

```
event log show -node impaired-node-name -messagename scsiblade.*
```

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

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

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

Shutting down the target controller

You shut down or take over the target controller using different procedures, depending on whether it is part of an HA pair or a stand-alone system.

Choices

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

- [Shutting down a stand-alone controller](#) on page 5

Shutting down a controller module in an HA pair

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

About this task

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

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

Steps

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

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

```
cluster show
```

The command produces output similar to the following:

```
Node  Health  Eligibility
-----
node1 true     true
node2 true     true
node3 true     true
node4 true     true
4 entries were displayed.
```

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

If...	Then...
All nodes show <code>true</code> for both health and eligibility.	Proceed to Step 3.
The impaired node shows <code>false</code> for health.	Proceed to the next step.
Any nodes show <code>false</code> for eligibility.	Resolve any cluster issues as needed before continuing with this procedure.
Any nodes other than the impaired node show <code>false</code> for health.	Correct the problems that cause the health issues on the nodes before continuing with this procedure.

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

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

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

If...	Then...
Neither node is in takeover mode	Go to the next step in this procedure.

If...	Then...
The healthy node took over the impaired node	The impaired node is in a state where you can begin removing it from the system chassis.
The impaired node took over the healthy node	<ol style="list-style-type: none"> a. Correct the problem that caused the takeover. b. Enter the cf giveback command (7-Mode) or storage failover giveback -fromnode <i>nodename</i> command (clustered Data ONTAP) from the impaired node console. c. Go back to the Step 1.

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

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

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

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

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

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

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

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

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

Shutting down a stand-alone controller

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


Steps

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

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

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

Attention: 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 located on the controller module to the right of the network ports, marked with a battery symbol. If the NVMEM LED is flashing, there is content in the NVMEM that has not been saved to disk. You need to reboot the controller module and proceed from the beginning of this procedure. If repeated attempts to cleanly shut down the controller module fail, be aware that you might lose any data that was not saved to disk.



1	NVMEM LED
----------	-----------

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

If your system uses...	Then...
AC power supplies	Unplug the power cords from the power source, and then remove the power cords.
DC power supplies	Remove the power at the DC source, and remove the DC wires, if necessary.

Removing the controller module and moving the components

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

About this task

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

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

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

Opening the system


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

Steps

1. If you are not already grounded, properly ground yourself.
2. Loosen the hook and loop strap binding the cables to the cable management 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.



4. Pull the cam handle downward and slide the controller module out of the system.

Moving the PCIe cards to the new controller module


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

Before you begin

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

Steps

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



1	Controller module side panel with thumbscrew
2	Controller module PCIe card


3. Remove the PCIe card from the controller module and set it aside.
Make sure that you keep track of which slot the PCIe card was in.
4. Repeat the preceding steps for the remaining PCIe cards in the old controller module.
5. Open the new controller module side panel, if necessary, slide off the PCIe card filler plate, as needed, and install the PCIe card.
Be sure that you properly align the card in the slot and exert even pressure on the card when seating it in the socket. The card must be fully and evenly seated in the slot.
6. Repeat the preceding step as needed for additional cards.
7. Close the side panel and tighten the thumbscrew.

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 cover
2	Boot device


2. Open the boot device cover and holding the device by its edges, gently lift it upward to remove it out of the housing.
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.
5. Check the boot device to make sure that it is seated squarely and completely in the socket or connector.
If necessary, remove the boot device and reseal it into the socket.
6. Close the boot device cover.

Moving the NVMEM battery

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


Steps

1. Open the CPU air duct.
2. 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	CPU air duct
2	NVMEM battery
3	NVMEM battery plug

3. Gently pull the tab on the battery housing, near the plug, away from the controller module side.
4. Place your forefinger at the far end of the battery housing and gently push it toward the CPU air duct.
You should see the tabs on the battery housing aligning with the notches in the controller module sheet metal.



1	NVMEM battery
2	Battery tabs
3	Notch on chassis with alignment arrow

5. Gently pull the battery housing toward the center of the controller module, and then lift the battery out of the controller module.
6. Align the tabs on the battery holder with the notches in the controller module side, and gently push the battery housing so that the notches are under the lip of the controller module side.
7. While gently pushing the battery against the sheet metal on the chassis to hold it in the battery guide, place the forefinger of your free hand against the battery housing behind the locking tab on the battery, and then gently push the battery housing away from the CPU air duct.
If it is properly aligned, the battery snaps into place on the side of the controller module. If it does not, repeat these steps.
8. In the new controller module, seat the battery in the holder and plug the battery cable into the socket.
The plug should lock down onto the socket on the controller module.

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. Open the CPU air duct.

3. Locate the DIMMs.


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



<p>1</p>	<p>NVMEM DIMMs 1 and 2</p> <p>Note: See <i>Replacing an NVMEM battery and NVMEM DIMMs in a 32xx system</i> for information about removing these two DIMMs.</p>
<p>2</p>	<p>System DIMMs 1 through 4</p> <p>The number of DIMMs in your system will vary:</p> <ul style="list-style-type: none"> • In the 3210 and 3240 models, only DIMM sockets 1 and 2 are populated. • In all other 32xx models, all DIMM sockets are populated.
<p>3</p>	<p>DIMM sockets</p> <p>The NVMEM DIMM sockets have white DIMM locking tabs, while the system DIMM sockets have black locking tabs.</p>

4. 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.

5. 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.



Caution: The DIMMs are located very close to the CPU heat sync, which might still be hot. Avoid touching the CPU heat sync when removing the DIMM.

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

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

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

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.

Important: Make sure that you install the NVMEM DIMMs only in the NVMEM DIMM slots.

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 needed.

Installing the new controller module and booting the system

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

About this task

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

Steps

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

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

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

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

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

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

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

Verifying and setting the HA state of the controller module

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

Steps

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

```
ha-config show
```

The HA state should be the same for all components.

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

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

```
ha-config modify controller [ha | non-ha]
```

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

3. If the displayed system state of the chassis does not match your system configuration, set the HA state for the chassis by entering the following command:

```
ha-config modify chassis [ha | non-ha]
```

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

Restoring the FC configuration for HA pairs

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

Before you begin

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

Steps

1. From the healthy node, verify the values of the FC configuration on the replacement node by entering the applicable command:

If your system is operating in...	Then issue this command...
7-Mode or Data ONTAP versions prior to 8.0	<code>partner fcadmin config</code>
Clustered Data ONTAP	<code>system node run -node <i>healthy-node-name</i> partner fcadmin config</code>

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

If the FC variables are...	Then...
The same as you recorded earlier	Proceed to the next step in this procedure.

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

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

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

About this task

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

- The *impaired node* is the node on which you are performing maintenance.
- 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 impaired node.

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

Steps

1. If you have not already done so, halt the replacement node to display the LOADER prompt.
2. On the healthy node, check the system time:
`date`
3. At the LOADER prompt, check the date and time on the replacement node:

```
show date
```

The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node:

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

5. If necessary, set the time in GMT on the replacement node:

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

6. At the LOADER prompt, confirm the date and time on the replacement node:

```
show date
```

The date and time are given in GMT.

Installing the firmware after replacing the controller module

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

About this task

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

Steps

1. Check the configuration of the SP by entering the following command at the LOADER prompt:

```
sp status
```

Note the version of the SP firmware and cross-check it and update it, if needed, when you log into the NetApp Support Site in the following steps.

2. Log into the SP from an administration host to install the firmware, using the following command:

```
ssh username@SP_IP_address
```

For more information about accessing the SP, see the *Data ONTAP System Administration Guide for Cluster-Mode*.

3. Log in to the NetApp Support Site, select the most current version of firmware for your system from those listed at support.netapp.com/NOW/cgi-bin/fw, and then follow the instructions for downloading and installing the new firmware.

Note: You can also take this opportunity to download and install the SP firmware and BIOS on the healthy node, if needed.

Running diagnostics tests (controller replacement)

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

Before you begin

- Your system must be at the LOADER prompt to start System Level Diagnostics.
- For Data ONTAP 8.1.x and earlier, you need loopback plugs to run tests on storage interfaces. You can test these components without the loopback plugs, but should verify from the long test output that the only failures that occurred were caused by the lack of loopback plugs.

About this task

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

Steps

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

```
boot_diags
```

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

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

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

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

The Maintenance mode prompt (`*>`) appears.

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

```
sldiag device clearstatus
```

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

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

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

- `bootmedia` is the system booting device.
- `env` is motherboard environmentals.
- `mem` is system memory.
- `nic` is a Network Interface Card.
- `nvmem` is a hybrid of NVRAM and system memory.
- `sas` is a Serial Attached SCSI device not connected to a disk shelf.
- `serviceproc` is the Service Processor.

Note the available devices on the system.

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

If you want to run...	Then use this procedure...
Several diagnostic tests concurrently	Complete the following substeps: on page 17
A single device diagnostic test at a time	Complete the following substeps: on page 18

Running diagnostics tests concurrently (controller replacement)

You can choose to run diagnostics tests concurrently if you want a single organized log of all the test results for all the devices.

About this task

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

Steps

1. Review the enabled and disabled devices in the output from Step 4 in the preceding procedure and determine which ones you want to run concurrently.
2. List the individual tests for the device:
`sldiag test show -dev dev_name`
3. Verify that the tests were modified:
`sldiag test show`
4. Repeat steps 1 through 5 for each device.
5. Run diagnostics on all the devices:
`sldiag device run`

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

The test is complete when the following message is displayed:

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

6. After the test is complete, verify that there are no hardware problems on your storage system:
`sldiag device status -long -state failed`
7. Correct any issues found and repeat this procedure.

Result

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

Running diagnostics tests individually (controller replacement)

You can run diagnostics tests individually if you want a separate log of all the test results for each device.

Steps

1. Clear the status logs by entering the following command:
`sldiag device clearstatus`
2. Display the available tests for the selected devices:

Device type	Command
boot media	<code>sldiag test show -dev boot media</code>
env	<code>sldiag test show -dev env</code>
mem	<code>sldiag test show -dev mem</code>
nic	<code>sldiag test show -dev nic</code>
nvmem	<code>sldiag test show -dev nvmem</code>
sas	<code>sldiag test show -dev sas</code>

The available device tests are displayed.

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

```
sldiag test 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 do want to run:

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

5. Run the selected tests:

Device type	Command
boot media	sldiag test run -dev boot media
env	sldiag test run -dev env
mem	sldiag test run -dev mem
nic	sldiag test run -dev nic
nvmmem	sldiag test run -dev nvmmem
sas	sldiag test 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 test status -dev boot media -long -state failed
env	sldiag test status -dev env -long -state failed
mem	sldiag test status -dev mem -long -state failed
nic	sldiag test status -dev nic -long -state failed
nvmmem	sldiag test status -dev nvmmem
sas	sldiag test 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	<p>a. Clear the status logs by entering the following command:</p> <pre data-bbox="456 306 818 331">sldiag device clearstatus</pre> <p>b. Verify that the log is cleared by entering the following command:</p> <pre data-bbox="456 396 747 422">sldiag device status</pre> <p>The following default response is displayed:</p> <pre data-bbox="456 457 971 483">SLDIAG: No log messages are present.</pre> <p>You have completed system-level diagnostics.</p>
Resulted in some test failures	<p>Determine the cause of the problem:</p> <p>a. Exit Maintenance mode by entering the following command:</p> <pre data-bbox="456 644 516 669">halt</pre> <p>After you issue the command, wait until the system stops at the LOADER prompt.</p> <p>b. Turn off or leave on the power supplies, depending on how many controller modules are in the chassis:</p> <ul data-bbox="456 764 1507 879" style="list-style-type: none"> • If you have two controller modules in the chassis, leave the power supplies turned on to provide power to the other controller module. • If you have one controller module in the chassis, turn off the power supplies and unplug them from the power sources. <p>c. Check the controller module you are servicing and verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system.</p> <p>d. Boot the controller module you are servicing, interrupting the boot by pressing Ctrl-C when prompted. This takes you to the Boot Menu:</p> <ul data-bbox="456 1077 1511 1192" style="list-style-type: none"> • If you have two controller modules in the chassis, fully seat the controller module you are servicing in the chassis. The controller module boots up when fully seated. • If you have one controller module in the chassis, connect the power supplies and turn them on. <p>e. Select Boot to maintenance mode from the menu.</p> <p>f. Exit Maintenance mode by entering the following command:</p> <pre data-bbox="456 1304 516 1329">halt</pre> <p>After you issue the command, wait until the system stops at the LOADER prompt.</p> <p>g. Enter boot_diags at the prompt and rerun the system-level diagnostic test.</p>

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

Recabling the system

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

Step

1. Recable the system, as needed.

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

Reassigning disks

On systems running Data ONTAP 8.1.x and earlier, you must reassign the disks attached to the target node (the node on which the controller module has been replaced) so that they point to the new system ID. This is because the new controller module has a new system ID.

Choices

- [Reassigning disks on a system operating in 7-Mode](#) on page 21
- [Reassigning disks on a system operating in clustered Data ONTAP](#) on page 23

Reassigning disks on a system operating in 7-Mode

You must reassign disks before you boot the software. Some of the steps are different depending on whether the system is stand-alone or in an HA pair.

About this task

- You must apply the commands in these steps on the correct systems:
 - The *impaired node* is the node on which you are performing maintenance.
 - The *healthy node* is the HA partner of the impaired node.
- Do not issue any commands relating to aggregates until the entire procedure is completed.
- If your system has 500 or more disk drives, the version of Data ONTAP you are running determines whether you can reassign the disks:

If you are running this version of Data ONTAP...	Then...
<ul style="list-style-type: none">◦ 8.0.3 or later in the 8.0 release family◦ 8.1.1 or later in the 8.1 release family	The operation is supported but you must contact technical support for assistance.
<ul style="list-style-type: none">◦ Any release in the 7.3 release family◦ 8.0.2 or earlier in the 8.0 release family◦ 8.1 in the 8.1 release family	You cannot reassign more than 500 disks from one controller to another by using the <code>disk reassign</code> command. If you try to do so, the system reports an error and you must contact technical support.

Steps

1. If you have not already done so, reboot the impaired 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 new system IDs by entering the following command:

```
disk show -v
```

Note: Make note of the new system ID, which is displayed in the Local System ID field.

Example

The example contains the following information:

- *system-2* is the impaired node, which is undergoing maintenance.
- *system-1* is the healthy node.
- The new system ID is *118065481*.
- The old system ID is *118073209*, which is still assigned to the disks owned by *system-2*.

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

  DISK   OWNER                POOL   SERIAL NUMBER   HOME
-----
0b.29   system-1 (118065578) Pool0  J8XJE9LC        system-1 (118065578)
0a.27   system-2 (118073209) Pool0  J8Y478RC        system-2 (118073209)
.
.
.
```

- If the controller module is in an HA pair, on the healthy node, enter the following command to ensure that any coredumps on the impaired node are saved:

```
partner savecore
```

- Reassign disk ownership based on your system's configuration:

If the controller module is... **Then perform these steps on the applicable node...**

In an HA pair

- Halt the impaired node by entering the following command on the impaired node:
halt
- Confirm that the impaired node has been taken over by entering the following command on the healthy node:
cf status
If the impaired node has not been taken over, you should use the `cf takeover` command on the healthy node to take it over.
- On the healthy node, enter the following command to enter advanced privilege mode:
priv set advanced
- On the healthy node, reassign disk ownership (for FAS systems) or LUN ownership (for V-Series systems), by using the system ID information obtained from the `disk show -v` command:
disk reassign -s *old system ID* -d *new system ID*
Continuing the example seen in Step 2:
 - The *old system ID* is *118073209*.
 - The *new system ID* is *118065481*.

Stand-alone

- Reassign disk ownership by entering the following command at the Maintenance mode prompt of the *impaired* node:
disk reassign -s *old system ID* -d *new system ID*
Continuing the example seen in Step 2:
 - The *old system ID* is *118073209*.
 - The *new system ID* is *118065481*.
-

- Verify that the disks (or V-Series LUNs) were assigned correctly by entering the following command:

```
disk show -v
```

Note: If your system is in an HA pair, you must run this command on the healthy node.

Example

Make sure that the disks belonging to the impaired node show the new system ID for the impaired node. In the following example, the disks owned by system-2 now show the new system ID, 118065481:

```
system-1> disk show -v
  DISK   OWNER                POOL   SERIAL NUMBER   HOME
-----
0b.17   system-2 (118065481) Pool0  J8Y0TDZC        system-2 (118065481)
```

```
0a.17 system-1 (118065578) Pool0 J8Y09DXC system-1 (118065578)
.
.
.
```

6. If you have a stand-alone system, exit Maintenance mode by entering the following command on the impaired node:
halt
7. After the impaired node displays the LOADER prompt, enter the following command to boot the operating system:
boot_ontap
For a system in an HA pair, this command puts the node in Waiting for Giveback state.
8. If the system is in an HA pair, do a giveback and confirm that the HA pair is healthy:
 - a. On the healthy node, enter the following command to return to standard privilege mode:
priv set admin
 - b. On the healthy node, return storage to the impaired node by entering the following command:
cf giveback
 - c. After the giveback operation is complete, enter the following command to check that the HA pair is healthy and takeover is possible:
cf show

Reassigning disks on a system operating in clustered Data ONTAP

You must assign the NVRAM system ID of the new controller module to the disks in the system.

About this task

- It is important that you apply the commands in these steps on the correct systems:
 - The *impaired node* is the node on which you are performing maintenance.
 - The *healthy node* is the HA partner of the impaired node.
- The impaired node must have been taken over by its partner.
If the impaired node has not been taken over, you should use the `cf takeover` command on the healthy node to take it over.
- If your system has 500 or more disk drives, the version of Data ONTAP you are running determines whether you can reassign the disks:

If you are running this version of Data ONTAP...	Then...
<ul style="list-style-type: none"> ◦ 8.0.3 or later in the 8.0 release family ◦ 8.1.1 or later in the 8.1 release family 	The operation is supported but you must contact technical support for assistance.
<ul style="list-style-type: none"> ◦ 8.0.2 or earlier in the 8.0 release family ◦ 8.1 in the 8.1 release family 	You cannot reassign more than 500 disks from one controller to another by using the <code>disk reassign</code> command. If you try to do so, the system reports an error and you must contact technical support.

Steps

1. Complete the following substeps on the healthy node:
 - a. Log in as admin and enter the password.
 - b. Enter the following command:
run local
 - c. Enter the following command to obtain the system ID of the failed NVRAM adapter:

disk show

Example

The command displays system and disk information, as shown in the following example. The old system ID of impaired node, 101174200, appears to the right of the column labeled HOME:

```
healthy_node> disk show
DISK   OWNER          (103668010)  HOME          (101174200)  POOL
-----
0b.18  healthy_node   (103668010)  impaired_node (101174200)  Pool0
[...]
0b.22  healthy_node   (103668010)  impaired_node (101174200)  Pool0
[...]
0b.20  healthy_node   (103668010)  impaired_node (101174200)  Pool0
```

- d. Write down the system ID.

This is the old system ID, which you use later in this procedure.

- 2. Complete the following substeps on the impaired node:

- a. If you have not already done so, boot the system into Maintenance mode.
- b. Enter the following command to obtain the system ID of the new NVRAM adapter:

disk show

Example

```
*> disk show
disk show
Local System ID: 0101166306
```

- c. Write down the system ID.

This is the new system ID, which you use later in this procedure.

- d. Exit Maintenance mode and display the LOADER prompt by entering the following command:

halt

- e. On the impaired node, enter the following command at the LOADER prompt to set the `bootarg.mgwd.autoconf.disable` variable to `true` to disable auto configuration:

setenv bootarg.mgwd.autoconf.disable true

- f. On the impaired node, enter the following command at the LOADER prompt to ensure that the new controller module boots in clustered Data ONTAP:

setenv bootarg.init.boot_clustered true

Note: Do not issue any commands related to aggregates until the entire procedure is completed.

- 3. Complete the following substeps on the healthy node:

- a. On the healthy node, enter the following command to ensure that any coredumps on the impaired node are saved:

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

- b. On the healthy node node, enter the following command to enter advanced privilege mode:

priv set advanced

- c. Enter the following command to assign the system ID of the new NVRAM adapter to the disks:

disk reassign -s old_system_ID -d new_system_id

Attention: Be sure to perform this step correctly. If disks are not assigned correctly, the nodes will panic.

`old_system_ID` represents the system ID you recorded in step 1d.

new_system_id represents the system ID you recorded in Step 2c.

- d. Enter **y** at the following prompt:

Example

```
disk reassign -s 101174200 -o impaired_node -d 0101166306
Disk ownership will be updated on all disks previously belonging to
Filer with sysid 101174200.
Would you like to continue (y/n)? y
```

- e. Enter the following command to ensure that the system IDs have been correctly reassigned to the disks:

```
disk show
```

- f. Enter the following command to return to the admin privilege mode:

```
priv set admin
```

4. Complete the following substeps on the impaired node (the one containing the new NVRAM adapter):

- a. Restart the node by entering the following command:

```
autoboot
```

- b. As the software starts, access the boot menu by pressing Ctrl-C.

- c. Update and synchronize the flash-based configuration by selecting the option `Update flash from backup config` from the displayed menu.

The impaired node boots and displays the
Waiting for giveback...
message.

- d. Issue Ctrl-C to let the system complete booting up without a giveback.

It will then reboot a second time to implement the configuration from backup.

- e. Boot the Data ONTAP software by entering the following command:

```
autoboot
```

- f. Enter **y** when the startup process prompts you to confirm the system ID mismatch.

Example

```
WARNING: System id mismatch. This usually occurs when moving CF cards!
Override system id (y|n) ? [n] y
```

This puts the node in Waiting for Giveback state.

5. On the healthy node, return the storage to the node containing the new NVRAM adapter by entering the following command:

```
storage failover giveback -fromnode local
```

6. After the giveback operation is complete, enter the following command to check that the HA pair is healthy and takeover is possible:

```
storage failover show
```

Performing a final takeover and giveback from the impaired node

To ensure that the disk reassignment is successful on a system running certain versions of Data ONTAP 8.0.x or Data ONTAP 8.1.x, you must perform a final takeover and giveback from the impaired node.

About this task

You can skip these steps if you are running one of these versions of Data ONTAP:

- Data ONTAP 8.0.5 or later in the Data ONTAP 8.0 release family
- Data ONTAP 8.1.3 or later in the Data ONTAP 8.1 release family

It is important that you apply the commands in the steps in the exact order they are presented and on the correct system:

- The *impaired node* is the node on which you are performing maintenance.
- The *healthy node* is the HA partner of the impaired node.

Steps

1. Take over the healthy node by entering one of the following commands from the console of the impaired node:

If your system is configured in...	Then issue this command...
7-Mode	<code>cf takeover</code>
Clustered Data ONTAP	<code>storage failover takeover -bynode node</code>

2. Return control to the healthy node by entering one of the following commands from the console of the impaired node:

If your system is configured in...	Then issue this command...
7-Mode	<code>cf giveback</code>
Clustered Data ONTAP	<code>storage failover giveback -fromnode nodename</code>

A successful giveback ends with a message on the healthy node indicating successful giveback.

Note: If `waiting for Giveback` is not displayed prior to giveback, reboot the controller module. If this continues, contact technical support.

Restoring licenses on the replacement node

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

Choices

- [Installing licenses for the replacement node in 7-Mode](#) on page 27
- [Installing licenses for the replacement node in clustered Data ONTAP](#) on page 27

Installing licenses for the replacement node in 7-Mode

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

About this task

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

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

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

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.

If the site does not have the license keys you need, contact your sales or support representative.

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

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

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

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

4. Issue the following command to verify that the licenses have been installed:

```
license show
```

Installing licenses for the replacement node in clustered Data ONTAP

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

About this task

Until you install license keys, features requiring standard licenses will continue to be available to the replacement node. However, if the impaired node was the only node in the cluster with a license for the feature, no configuration changes to the feature are allowed. Also, using unlicensed features on the node might put you out of compliance with your license agreement, so you should install the replacement license key or keys on the replacement node as soon as possible. For detailed information about licensing, see the knowledgebase article [Data ONTAP 8.2 Licensing Overview and References](#) on the NetApp Support Site and the *Data ONTAP System Administration Guide for Cluster-Mode*.

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

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.

If the site does not have the license keys you need, contact your sales or support representative.

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

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

Restoring Storage Encryption functionality after controller module replacement

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

Steps

1. Enter the following command at the storage system prompt:

```
key_manager setup
```

2. Complete the steps in the wizard to configure Storage Encryption.

Ensure that a new passphrase is generated and that you select **Yes** to lock all drives.

3. Repeat Steps 1 and 2 on the partner node.

Do 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 by entering the following command:

```
disk encrypt show
```

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

5. On each node, load all authentication keys by entering the following command:

```
key_manager restore -all
```

6. On each node, verify that all keys are stored on their key management servers by entering the following command:

```
key_manager query
```

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

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 support.netapp.com, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Disposing of batteries

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

Related information

[Warranty Agreement, Safety Information, and Regulatory Notices at support.netapp.com](http://support.netapp.com)

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and www.ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Copyright information

Copyright © 1994–2014 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Customer Fitness, Cryptainer, CryptoShred, CyberSnap, Data Center Fitness, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, ExpressPod, FAServer, FastStak, FilerView, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, Snap Creator, SnapDirector, SnapDrive, SnapFilter, SnapIntegrator, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, VelocityStak, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by email to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

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