



Replacing PCIe cards and risers in a 62xx system

To replace a PCIe card or riser in the system, you must perform a specific sequence of tasks.

About this task

- You can use this procedure with all versions of Data ONTAP supported by your system.
In this procedure, a *Cluster-Mode system* refers to a system running Data ONTAP 8.x Cluster-Mode. A *7-Mode system* refers to a system running Data ONTAP 8.x 7-Mode or releases prior to Data ONTAP 8.0.
- This procedure refers to *HA pairs*, which in releases prior to Data ONTAP 8.0 were called *active/active configurations*.
- All other components in the system must be functioning properly; if not, you must contact technical support.

Steps

1. [Shutting down the node](#) on page 1
2. [Opening the system](#) on page 3
3. [Removing a PCIe card](#) on page 3
4. [Installing a PCIe card](#) on page 6
5. [Removing a riser](#) on page 6
6. [Installing a riser](#) on page 8
7. [Reinstalling the controller module and booting the system](#) on page 8
8. [Running diagnostics on a PCIe card](#) on page 9
9. [Configuring newly added Ethernet ports for cluster or management use \(Data ONTAP Cluster-Mode\)](#) on page 11
10. [Completing the replacement process](#) on page 11

Shutting down the node

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

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.

About this task

As part of this procedure, you leave the power supplies 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 partner node.
- If you have one controller module in the chassis, but is part of an HA pair, you should turn off the power supplies in the target node chassis.

Steps

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

If your system is configured in...	Then issue this command...
7-Mode	cf status
Cluster-Mode	storage failover show

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

If...	Then...
Neither node is in takeover mode	Go to the next step in this procedure.
The partner node took over the target node	The target node is in a state where you can begin removing it from the system chassis.
The target node took over the partner node	<ol style="list-style-type: none"> Correct the problem that caused the takeover. Run the <code>cf giveback</code> command (if in a 7-Mode system) or <code>storage failover giveback -fromnode nodename</code> command (if in a Cluster-Mode system) from the target node console. Go back to the beginning of this procedure.

3. Take over the target node by entering one of the following commands from the partner node's console:

If your system is configured in...	Then issue this command...
7-Mode	cf takeover
Cluster-Mode	storage failover takeover -bynode node

4. Take one of the following actions depending on your system configuration:

If your system has...	Then...
Two controller modules in the chassis	Do not shut off the power supplies.
One controller module in the chassis	Turn off the power supplies and unplug them from the power sources.

Shutting down the node in a stand-alone system

For a node that is in a stand-alone configuration, you must perform a clean shutdown (ensuring that all data has been written to disk) and disconnect the power supplies.

Steps

1. Enter one of the following commands from the system console:

If your system is configured in...	Then issue this command...
7-Mode	halt
Cluster-Mode	halt local

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

Opening the system

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

Steps

1. If you are not already grounded, properly ground yourself.
2. Unplug the system cables from the controller module, as needed, and keep track of where the cables were connected.
Leave the cables in the cable management tray so that when you reinstall the cable management tray, the cables are organized.
3. Grasp the cable management tray by the side, gently push one side of the tray so that the arm slides across and off the retaining pin, slide the other arm off the retaining pin on the other side of the controller module, and then lift it off the back of the controller module and set it aside.
4. Loosen the thumbscrew on the cam handle.
5. Push in the release latch on the left side of the controller module and slide the controller module out of the system and set it on a stable, antistatic surface.

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


Removing a PCIe card

To remove a PCIe card from the system, you must perform a specific sequence of steps.

Steps

1. Loosen the thumbscrew on the controller module side panel.

2. Swing the side panel open until it comes off the controller module.



1	Side panel with thumbscrew
2	PCIe card
3	PCIe riser card

3. Remove the PCIe card from the controller module and set it aside.
4. Repeat these steps as needed.
5. If you have a SAS or FC card in slot 1 of the controller module, remove the card from the controller module by completing the following substeps:

Installing a PCIe card

To install a PCIe card in the system, you must perform a specific sequence of steps.

Steps

1. Open the controller module side panel, if necessary, and slide off the PCIe card filler plate, as needed.
2. 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 adapter must be fully and evenly seated in the slot.
3. Repeat these steps as needed for additional PCIe cards.
4. If you are replacing a SAS or FC card in slot 1 of the controller module, complete the following substeps:
 - a. If necessary, loosen the two thumbscrews on the NVRAM and vertical I/O cover and remove the cover.
 - b. Install the card in slot 1 by aligning the card connector with the socket on the controller module and gently and evenly pushing down on the edges of the card to seat it in the socket.
 - c. Reinstall the NVRAM and vertical I/O cover and tighten the thumbscrews.
5. Close the side panel and tighten the thumbscrew.

Removing a riser


To remove the riser from the riser card from the controller module, you must perform a specific sequence of steps..

Before you begin

You must have removed all PCIe cards from the riser.

Steps

1. Using the FRU map on the controller module CPU air duct for reference, loosen the riser thumbscrew located at the top and middle of the riser card.
2. Pivot the riser arms up, and then pull the riser up and out of the sockets by the tabs on the riser arms.



1	Side panel with thumbscrew
2	PCIe card
3	PCIe riser card

Installing a riser

To install the riser in the controller, you must perform a specific sequence of steps.

Steps

1. While grounded, align the riser with the guide slots on the chassis, push the riser into the slots using the cam handles, and then rotate the cam handles to the closed position to seat the riser into the connector on the controller module.
2. Examine the riser to make sure that it is seated squarely and completely in the sockets.

If the riser is not seated squarely, remove and reinstall it.

3. Tighten the riser thumbscrew.
4. If necessary, reinstall the PCIe cards.

Reinstalling the controller module and booting the system

After you replace a component within the controller module, you must reinstall the controller module in the system chassis and boot it to a state where you can run diagnostic tests on the replaced component.

About this task

Note: 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.
2. Reinstall the cable management arms and recable the controller module, as needed.
When recabling, remember to reinstall the media converters (SFPs) if you are using fiber cables.
3. Reinstall the controller module:
 - a. Push the controller module all the way into the system.
 - b. Push the cam handle to finish seating the controller module.
 - c. Close the cam and tighten the thumbscrew.
 - d. Halt the boot process on the controller module undergoing maintenance:
 - For systems with two controller modules in the chassis in an HA pair, the system automatically begins to boot when the controller module is fully seated.
Press Ctrl-c to interrupt the boot process.
 - For stand-alone systems or systems in an HA pair with one controller module in the chassis, 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.

The boot process is halted at either the Boot menu or the Loader prompt.

Running diagnostics on a PCIe card

After installing a PCIe card, you should run diagnostics.

Before you begin

Before running diagnostics tests on your PCIe card you need to have it cabled according to the following guidelines:

- NIC interfaces must be cabled in a pair-wise manner, with adjacent ports connected together or connected through a switch. Ports e0M and e0P must not be connected together because of the internal switch connection.
- SAS card ports must be connected to storage or connected to the adjacent SAS ports. If the SAS cards or systems with more than two ports are not connected to storage, they must be connected to ports A to B, C to D, and so forth.
- The FC card ports must be connected to storage or the ports must be terminated with loopback plugs.

Steps

1. Complete the applicable step, depending on where the node halted during the boot process.

If the node halted at the...	Then...
Loader prompt	Continue with the procedure.
Boot menu	<ol style="list-style-type: none">a. Select the Maintenance mode option from the displayed menu.b. Enter the following command at the prompt: haltc. Continue with the procedure.

2. Enter the following command at the Loader prompt:

boot_diags

Note: You must run 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.

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

3. Enter the following command at the Maintenance mode prompt:

sldiag

For details about the `sldiag` command, see the `sldiag` man page.

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

sldiag device clearstatus

5. Verify that the log is cleared by entering the following command:

sldiag device status

The following default response is displayed:

```
SLDIAG: No log messages are present.
```

6. View the types of devices available for checking by entering the following command:

sldiag device show

7. Run the test for the type of PCIe card you installed by entering the following command at the Maintenance mode prompt:

sldiag device run -dev fcal/sas/ata/nic/toe

See the *System-Level Diagnostics Guide* for a description of other types of tests you can run.

8. View the status of the test by entering the following command:

```
sldiag device status
```

Your storage system provides the following output while the tests are still running:

```
There are still test(s) being processed.
```

After all the tests are complete, the following response appears by default:

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

9. Verify that no hardware problems resulted from the addition or replacement of hardware components on your system by entering the following command:

```
sldiag device status [-dev devtype] [-name device] -long -state failed
```

System-level diagnostics returns you to the prompt if there are no test failures, or lists the full status of failures resulting from testing the component.

10. Proceed based on the result of the preceding step.

If the system-level diagnostics tests...	Then...
Were completed without any failures	<ol style="list-style-type: none">a. Clear the status logs by entering the following command: <pre>sldiag device clearstatus</pre>b. Verify that the log is cleared by entering the following command: <pre>sldiag device status</pre><p>The following default response is displayed:</p><pre>SLDIAG: No log messages are present.</pre>c. Exit Maintenance mode by entering the following command: <pre>halt</pre>d. Enter the following command at the firmware prompt to reboot the storage system: <pre>boot</pre>e. If your system is in an HA pair, run the <code>cf giveback</code> command (if in a 7-Mode system) or <code>storage failover giveback</code> command (if in a Cluster-Mode system) from the partner node console. <p>You have completed system-level diagnostics.</p>
Resulted in some test failures	<p>Determine the cause of the problem.</p> <ol style="list-style-type: none">a. Exit Maintenance mode by entering the following command: <pre>halt</pre>b. Perform a clean shutdown and disconnect the power supplies.c. Verify that you have observed all the considerations identified for running system-level diagnostics, that cables are securely connected, and that hardware components are properly installed in the storage system.d. Reconnect the power supplies and power on the storage system.e. Rerun the system-level diagnostics tests.

Related information

System-Level Diagnostics Guide: now.netapp.com/NOW/knowledge/docs/hardware/NetApp/diag/sldiag.pdf

Configuring newly added Ethernet ports for cluster or management use (Data ONTAP Cluster-Mode)

When you install an Ethernet card in a Data ONTAP Cluster-Mode system, the ports are automatically configured as data ports. To use the ports as cluster ports or management ports, you must manually reconfigure them.

Step

1. Enter the following command to change the port configuration:

```
network port modify -node node -port port-name -role port-role
```

port-role can be *mgmt*, *cluster*, or *data*.

For more information, see the *Data ONTAP Cluster-Mode Network and File Access Management Guide* on the NetApp Support Site.

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, 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, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, 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 U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. 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 U.S.A. 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.

