



Replacing the I/O expansion module in a 62xx system

Replacing the I/O expansion module involves shutting down the system in which the I/O expansion module is installed and moving all cards and adapters from the old I/O expansion module to the new I/O expansion module.

Before you begin

All other components in the system must be functioning properly; if not, you must contact technical support.

About this task

- You can only replace the I/O expansion module in systems that already have an I/O expansion module installed.
- You cannot install the I/O expansion module into a system that contains only a single controller module.
- You must shut down the system that contains the I/O expansion module before you replace it.
- 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*.

1. [Shutting down the node](#) on page 1
2. [Removing the I/O expansion module](#) on page 2
3. [Installing the I/O expansion module](#) on page 4
4. [Running diagnostics on the I/O expansion module](#) on page 5
5. [Completing the replacement process](#) on page 7

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.

Next topics

[Shutting down a node in an HA configuration](#) on page 1

[Shutting down the node in a stand-alone system](#) on page 2

Shutting down a node in an HA configuration

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 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	<code>cf status</code>
Cluster-Mode	<code>storage failover show</code>

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</code> command (if in a Cluster-Mode system) from the target node console. Go back to the beginning of this procedure.

- 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	<code>cf takeover</code>
Cluster-Mode	<code>storage failover takeover -bynode node</code>

- If you are not already grounded, properly ground yourself.
- Turn off the power supplies and disconnect the power on the target node.

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

- Enter one of the following commands from the system console:

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

- 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.


Removing the I/O expansion module

To remove the I/O expansion module from the chassis, you must perform a specific sequence of steps.

Steps


- If you are not already grounded, properly ground yourself.
- Unplug the system cables from the I/O expansion module, as needed, and keep track of where the cables were connected.
Leave the cables in the cable management tray so that when you install the cable management tray in the new I/O expansion module, the cables will be organized.
- Loosen the thumbscrew on the cam handle.
- Pull the cam handle down and slide the I/O expansion module out of the chassis until it catches, press the release tab on the side of the module and then slide it out of the system, and then place it on an anti-static surface.

5. Remove the side panels from the I/O expansion module by loosening the side panel thumbscrew and swinging the side panel out of the way.



1	Right side panel, side closest to the the cam handle, of the I/O expansion module
2	Right PCIe cards There can be up to four PCIe cards per side of the I/O expansion module.
3	Right side PCIe riser card
4	Left side PCIe riser card
5	Left PCIe cards There can be four up to PCIe cards per side of the I/O expansion module.
6	Left side panel, side closest to the cam handle, of the I/O expansion module

6. Remove the PCIe cards and then remove the two riser cards, noting their locations so that you can install them into the same slots in the new I/O expansion module.
7. If your I/O expansion module has SAS or FC adapters, remove the cover in the middle of the I/O expansion module by loosening the thumbscrew and swinging the panel upward, and then remove the SAS or FC adapters.



1	SAS or FC adapter
2	SAS or FC adapter

Installing the I/O expansion module

To install the I/O expansion module, you must perform a specific sequence of steps.

Steps

1. Open the new I/O expansion module side panels, as needed.

2. Install the riser cards into the I/O expansion module by aligning the riser card with the slot on the I/O expansion module, and gently seating the riser card into the socket, and then lowering the cam handles into place and tightening the thumbscrew.
3. Install the individual PCIe cards into the proper I/O expansion module slots by aligning each PCIe card with the target slot, and then gently seating it into the socket.

Make sure that you install the PCIe cards in the same slots in the new I/O expansion module as where they were in the old I/O expansion module.

4. Install your SAS or FC cards into slots 11 and 12 in the center of the I/O expansion module, and then close and secure the cover.
5. Close the side panels and tighten the side panel thumbscrews.
6. Align the I/O expansion module with the bottom opening of the bottom bay of the target chassis (labeled B on the bay), and gently slide the I/O expansion module all the way into the chassis using the cam handle.
The I/O expansion module cam handle begins to engage prior to fully seating into the chassis. Use the cam handle to fully seat the module into the chassis.
7. Push the cam handle all the way to the closed position, and then tighten the thumbscrew on the cam handle.
8. Reinstall the cable management arms, and recable your I/O expansion module.

Running diagnostics on the I/O expansion module

After installing the I/O expansion module, you should run diagnostics on the individual components to ensure that they and the new I/O expansion module are working properly.

Before you begin

Before running diagnostics tests on the PCIe cards in your I/O expansion module, 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. Enter `sldiag device show` and then connect the ports to each other, based on the order in which the port names appear on the console. Do not connect ports e0M and e0P 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 2 ports are not connected to storage, connect ports A to B, C to D, and so forth.
- The FC card ports must be connected to storage or the ports terminated with loopback plugs.

Steps

1. Reconnect the system to the power source and turn on the power.
2. Halt the boot process by pressing Ctrl-c.
The boot process is halted at either the Boot menu or the Loader prompt.
3. 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: halt c. Continue with the procedure.

4. 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.

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

```
sldiag
```

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

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

```
sldiag device clearstatus
```

7. 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.
```

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

```
sldiag device show
```

9. 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.

10. 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>
```

11. 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.

12. 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"> Clear the status logs by entering the following command: sldiag device clearstatus Verify that the log is cleared by entering the following command: sldiag device status The following default response is displayed: <pre>SLDIAG: No log messages are present.</pre> Exit Maintenance mode by entering the following command: halt Enter the following command at the firmware prompt to reboot the storage system: boot If your system is in an HA pair, run the cf giveback command (if in a 7-Mode system) or storage failover giveback 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"> Exit Maintenance mode by entering the following command: halt Perform a clean shutdown and disconnect the power supplies. 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. Reconnect the power supplies and power on the storage system. Rerun the system-level diagnostics tests.

Related information

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

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 NetApp 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.

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: now.netapp.com/NOW/knowledge/docs/hardware/hardware_index.shtml

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, ApplianceWatch, ASUP, AutoSupport, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort,

FAServer, FilerView, FlexCache, FlexClone, 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), ONTAPI, OpenKey, RAID-DP, ReplicatorX, SANscreen, SecureAdmin, SecureShare, Select, Shadow Tape, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, 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, and Web Filer 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.