



Replacing the 10-GbE riser card in a 62xx system

To replace the 10-GbE riser card, you must shut down the target controller module, open the system, remove the old 10-GbE riser card and install the new one, close the system, and then reboot it.

About this task

The 10-GbE riser card is not hot-swappable.

Next topics

[Shutting down the node](#) on page 1

[Opening the system](#) on page 2

[Removing the on-board 10-GbE riser card](#) on page 3

[Installing the on-board 10-GbE riser card](#) on page 4

[Reinstalling the controller module and booting the system](#) on page 5

[Running diagnostics on the on-board 10-GbE riser card](#) on page 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 pair](#) on page 1

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

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 | <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">a. Correct the problem that caused the takeover.b. 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.c. 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 | <code>cf takeover</code> |
| Cluster-Mode | <code>storage failover takeover -bynode node</code> |

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 | <code>halt</code> |
| Cluster-Mode | <code>halt local</code> |

2. If you are not already grounded, properly ground yourself.
3. 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 in the controller module, the cables will be 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. Pull the cam handle downward and slide the controller module out of the system until it catches.
6. 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 the on-board 10-GbE riser card

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

Steps

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



1

10-GbE riser card

2. Pull the card straight up and lift it out of the controller module, and then set it aside.

Installing the on-board 10-GbE riser card

To install the 10-GbE riser card into the controller module, you must perform a specific sequence of steps.

Steps

1. Using the FRU label inside the controller module, locate the correct socket for the 10-GbE riser card.
The socket for the card is directly in-line with the PCIe riser card.
2. Loosen the thumbscrew on the replacement 10-GbE riser card and pivot the cam handle arms up, if necessary.

3. Align the 10-GbE riser card with the slot on the controller module and gently push it into the slot using the cam handles.
4. Pivot the cam handle arms down to the closed position and tighten the thumbscrew.

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.
 - For systems in an HA pair, press Ctrl-c.
 - For stand-alone systems, 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 the on-board 10-GbE riser card

After installing the 10-GbE riser card into the system, you should run diagnostics.

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. Run the test for the 10-GbE riser card you installed by entering the following command at the Maintenance mode prompt:

```
sldiag device run -dev nic
```

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

7. 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>
```

8. 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.

9. 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 <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"> 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.

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.