Replacing a chassis

To replace the chassis, you must remove the controller module or modules and I/O expansion module, when present, from the old chassis, remove the old chassis from the equipment rack or system cabinet, install the new chassis of the same model as the old chassis, and then reinstall the components into the new chassis.

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 use this procedure with all versions of Data ONTAP supported by your system.

- This procedure is written with the assumption that you are moving the controller module or modules and I/O expansion module to the new chassis, and that the chassis is a new component from NetApp.

Steps

Shutting down the nodes

You must shut down the node or nodes in the chassis prior to moving them to the new chassis.

About this task

If you cannot gracefully shut down the nodes due to system constraints such as backplane failure, call technical support for additional guidance.

Steps

1. Change to the advance privilege level:

   `set -privilege advanced`

   a. How you proceed depends on where epsilon is assigned:

      - If the target chassis does not contain the epsilon node in it, then do to the next step.

      - If the target chassis does contain the epsilon node in it, remove and

2. Display the status of the nodes:

   `cluster show -epsilon *`

You will see a screen similar to the following:

```
Node          Health  Eligibility   Epsilon
------------  -------  ------------  ------------
node1         true    true          true
node2         true    true          false
node3         true    true          false
node4         true    true          false

4 entries were displayed.
```

Note: Epsilon must not be on a node in the chassis to be replaced.

3. If epsilon is assigned to a node in the chassis you are replacing, reassign it to a different node in a chassis that is healthy:

   a. Remove epsilon from the node:

      cluster modify -node node1 -epsilon false

   b. Assign epsilon to a node in a healthy chassis:

      cluster modify -node node4 -epsilon true

   c. Display epsilon for the cluster:

      cluster show -epsilon *

   *Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators*

4. If your system has two controller modules, you must disable the HA pair with one of the following commands:

If your system is configured as a...	Then issue this command...
System operating in 7-Mode	**cf disable**
System running clustered Data ONTAP with two nodes in the cluster	**node::> cluster ha modify -configured false** **node::> storage failover modify -node node0 -enabled false**
System running clustered Data ONTAP with more than two nodes in the cluster	**node::> storage failover modify -node node0 -enabled false**

5. Halt the node:

If your system is configured as a...	Then issue this command...
System operating in 7-Mode	**halt**
System running clustered Data ONTAP	**system node halt -node <node_name>**

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.
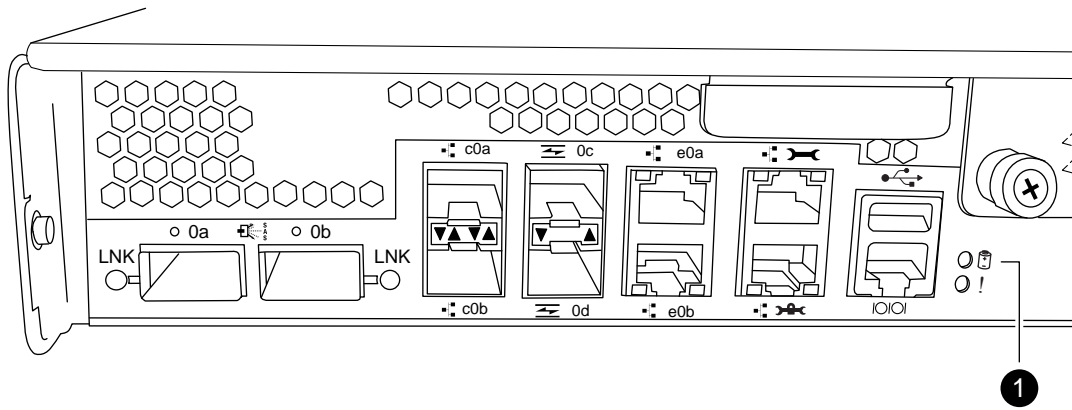
1	NVMEM LED

6. Repeat the preceding step for the second controller module, if present.

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

8. Turn off the power supplies and unplug both power cords from the power source.

Removing the controller module and/or I/O expansion module

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

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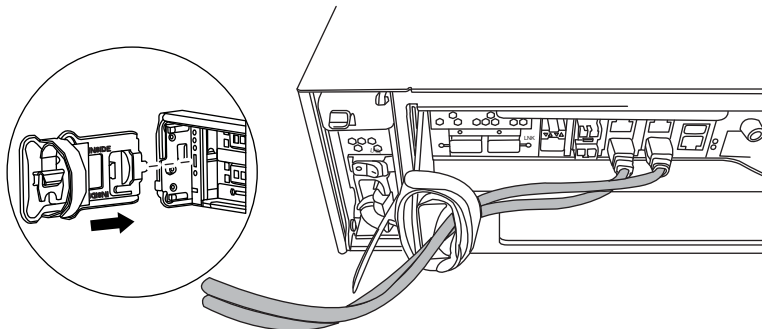
4. Pull the cam handle down and slide the controller module out of the system.

5. Repeat the preceding steps if you need to remove the second controller module, or if you need to remove the blanking panel.

Removing the chassis from the equipment rack or system cabinet

You must remove the existing chassis from the equipment rack or system cabinet before you can install the new one.

Steps

1. Remove the screws from the chassis mount points.

2. With the help of two or three people, slide the empty chassis off the rack rails and set it aside.

Installing the new chassis into the equipment rack or system cabinet

To install the new chassis in the equipment rack or system cabinet, you must perform a specific sequence of steps.

Steps

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

2. Using two or three people, install the empty chassis into the equipment rack or system cabinet by guiding the chassis onto the rack rails in a system cabinet or *L* brackets in an equipment rack.

3. Slide the chassis all the way into the equipment rack or system cabinet.

4. Secure the rear of the chassis to the support rail, if applicable, using the screws you removed from the defective chassis.

   Note: The rear of the 32xx chassis is not held down in a two-post rack. If you are installing the chassis in a system cabinet, you must use a metal hold-down strap and attach it to the system cabinet using the screws for the strap.

5. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the chassis.

Reinstalling the controller module and I/O expansion module and booting the system

To install and cable controller modules, and, if applicable, the I/O expansion module in the new chassis and then boot the system, you must perform a specific sequence of steps.

Steps

1. Repeat the previous steps if you need to install the I/O expansion module in the chassis

2. Reinstall the cable management device on the controller module and I/O expansion module, if present.

3. Recable the console port only.

4. The next step depends on your system configuration.

If your system is in...	Then...
A stand-alone configuration, single controller module only	Reinstall the blanking panel and then go to the next step.
An HA pair with a second controller module	Repeat the preceding steps for the second controller module.

If your system is in...	Then...
An HA pair with an I/O expansion module or stand-alone system with an I/O expansion module instead of a second controller module	**a.** Install the I/O expansion module using the preceding steps for installing the first controller module. **b.** Recable the adapters in the I/O expansion module.

5. Connect the power supplies to different power sources, and then turn them on.

6. Boot each node to Maintenance mode.

Verifying and setting the HA state of the chassis

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

Steps

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

   ha-config show

   The HA state should be the same for all components.

2. Proceed based on the results of Step 1.

If the displayed system state for the chassis...	Then...
Matches your system configuration	Go to the next step.
Does not match your system configuration	**a.** Set the HA state for the chassis: **ha-config modify chassis *HA-state*** The value for *HA-state* can be one of the following: • **ha** • **non-ha** • **mcc** **b.** Confirm that the setting has changed: **ha-config show**

3. If you have not already done so, recable the rest of your system.

4. The next step depends on your system configuration.

If your system is in...	Then...
A stand-alone configuration	**a.** Exit Maintenance mode: **halt** **b.** Go to *"Completing the replacement process* on page 8.

If your system is in...	Then...
An HA pair with a second controller module	Exit Maintenance mode: **halt** The LOADER prompt appears.
An HA pair with an I/O expansion module instead of a second controller module	Exit Maintenance mode by entering the following command: **halt** The LOADER prompt appears.

Running diagnostics on the chassis

After installing a new chassis, you should run interconnect diagnostics.

Steps

1. 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.

   You are prompted to continue the boot. Enter **y**, and then the Maintenance mode prompt (*>) appears.

2. Repeat the previous step on the second controller if you are in an HA configuration.

   > **Note:** Both controllers must be in Maintenance mode to run the interconnect test.

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

   sldiag device clearstatus

4. 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.

5. Enable the interconnect diagnostics tests by entering the following command from the Maintenance mode prompt:

   sldiag device modify -dev interconnect -sel enable

   > **Note:** The interconnect tests are disabled by default and must be enabled to run separately.

6. Run the interconnect diagnostics test by entering the following command from the Maintenance mode prompt:

   sldiag device run -dev interconnect

   > **Note:** You only need to run the interconnect test from one controller.

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 replacement of the chassis by entering the following command from the Maintenance mode prompt:

`sldiag device status -dev interconnect -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	**a.** Clear the status logs by entering the following command:
	
	`sldiag device clearstatus`
	
	b. 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>
	c. Exit Maintenance mode on both controllers by entering following command:
	
	`halt`
	
	    Attention: You must exit Maintenance mode on both controllers before proceeding any further.
	
	d. Enter the following command on both controllers at the LOADER prompt:
	
	`boot_ontap`
	
	e. If your system is a:
	
	  • 7-Mode system, enter
	
	    `cf enable`
	
	  • System running clustered Data ONTAP with two nodes in the cluster, enter the following commands:
	
	    `node::> cluster ha modify -configured true`
	
	    `node::> storage failover modify -node node0 -enabled true`
	
	  • System running clustered Data ONTAP with more than two nodes in the cluster, enter the following command:
	
	    `node::> storage failover modify -node node0 -enabled true`
	
	You have completed system-level diagnostics.

If the system-level diagnostics tests...	Then...
Resulted in some test failures	Determine the cause of the problem.

    a. Exit Maintenance mode by entering the following command:

       `halt`

    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

Completing the replacement process

After you replace the part, you can return the failed part to NetApp, as described in the RMA instructions shipped with the kit. Contact technical support at *NetApp Support*, 888-463-8277 (North America), 00-800-44-638277 (Europe), or +800-800-80-800 (Asia/Pacific) if you need the RMA number or additional help with the replacement procedure.

Copyright information

Trademark information

NetApp, the NetApp logo, Go Further, Faster, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, and WAFL are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. A current list of NetApp trademarks is available on the web at *http://www.netapp.com/us/legal/netapptmlist.aspx*.

Cisco and the Cisco logo are trademarks of Cisco in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

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