
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 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

1. [Shutting down the target node or nodes](#) on page 1
2. [Removing the controller module and/or I/O expansion module](#) on page 4
3. [Removing the chassis from the equipment rack or system cabinet](#) on page 5
4. [Installing the new chassis into the equipment rack or system cabinet](#) on page 5
5. [Reinstalling the controller module and I/O expansion module and booting the system](#) on page 5
6. [Verifying and setting the HA state of the chassis](#) on page 6
7. [Running diagnostics on the chassis](#) on page 7
8. [Completing the replacement process](#) on page 9

Shutting down the target node or nodes

You shut down or take over the target nodes using different procedures, depending on the storage system hardware configuration.

Shutting down nodes when replacing a chassis

You must shut down all of the nodes in the chassis before moving them to the new chassis.

About this task

If you cannot cleanly shut down the nodes because of system constraints, such as backplane failure, contact technical support for additional guidance.

Steps

1. Change to the advanced privilege level: `set -privilege advanced`
2. How you proceed depends on where epsilon is assigned:
 - If the target chassis does not contain the epsilon node, go to the step [3](#) on page 2

- If the target chassis contains the epsilon node, remove the node, and then reinstall the node.

3. Display the status of the nodes: `cluster show -epsilon *`

The following example shows the status of the nodes:

```

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 assigned to a node in the chassis that is to be replaced.

4. If epsilon is assigned to a node in the chassis that you are replacing, reassign epsilon to a 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 *`

ONTAP 9 System Administration Reference

5. If your system has two controller modules in the same chassis, disable the HA pair:

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

6. If your system is in a cluster and has a controller module and I/O expansion module in the same chassis, take over the node in the impaired chassis:

If your system is configured as a...	Then enter this command...
System operating in 7-Mode	<code>cf takeover</code>
System running clustered Data ONTAP	<code>storage failover takeover -ofnode <i>impaired_node_name</i></code>

7. Halt the node:

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

Attention: You must perform a clean system shutdown before replacing system components to avoid losing unwritten data in the NVRAM. The NVRAM adapter has a red LED that blinks if there is unwritten data in the NVRAM. If this LED is flashing red after you enter the `halt` command, you must reboot your system, and then try halting the node again.

8. Repeat step 7 on page 2 for the second controller module, if present.
9. If you are not already grounded, properly ground yourself.
10. Turn off the power supplies and unplug both power cords from the power source.

Shutting down a controller module in a two-node MetroCluster configuration

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

Steps

1. If the system is running 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. Use the `metrocluster check run`, `metrocluster check show` and `metrocluster check config-replication show` commands to make sure no configuration updates are in progress or pending.
3. If the impaired node has not switched over, perform the switchover operation from the healthy node:


```
metrocluster switchover
```
4. Monitor the completion of the switchover:


```
metrocluster operation show
```

Example

```
mcclA::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: in-progress
End time: -
Errors:

mcclA::*> metrocluster operation show
Operation: Switchover
Start time: 10/4/2012 19:04:13
State: successful
End time: 10/4/2012 19:04:22
Errors: -
```

5. Shut down the impaired node.

The method you use to shut down the node depends on whether remote management using a Service Processor (SP) is used:

Is the SP configured?	Then...
Yes	Log in to the impaired node's SP and issue the following command: system power off
No	At the impaired node's prompt, press Ctrl-C and respond y to halt the node.

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

7. Turn off the power supplies for the impaired node.

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 tray, 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 tray so that when you reinstall the cable management tray, the cables are organized.
3. Pull the cam handle down and slide the controller module out of the system.
4. 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.
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.

Note: If the system is in a system cabinet, you might need to remove the rear tie-down bracket.

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 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.
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	<ol style="list-style-type: none">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	<ol style="list-style-type: none"> a. Set the HA state for the chassis: <pre>ha-config modify chassis HA-state</pre> The value for <i>HA-state</i> can be one of the following: <ul style="list-style-type: none"> • ha • non-ha • mcc b. Confirm that the setting has changed: <pre>ha-config show</pre>

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	<ol style="list-style-type: none"> a. Exit Maintenance mode: <pre>halt</pre> b. Go to "Completing the replacement process" on page 9.
An HA pair with a second controller module	Exit Maintenance mode: <pre>halt</pre> 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: <pre>halt</pre> The LOADER prompt appears.

Running diagnostics on the chassis

After you have replaced the chassis in a controller-controller configuration, you should run interconnect diagnostics on the replacement chassis. If you have a controller-I/O expansion module configuration, you can skip this procedure.

Steps

1. Run the system-level diagnostics at the LOADER prompt: `boot_diags`

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

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

2. Repeat step 1 on page 7 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: `sldiag device clearstatus`

4. Verify that the log is cleared: `sldiag device status`

The following default response is displayed:

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

5. Enable the interconnect diagnostics at 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 at the Maintenance mode prompt: `sldiag device run -dev interconnect`

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

7. View the status of the test: `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 at the Maintenance mode prompt: `sldiag device status -dev interconnect -long -state failed`

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

9. Proceed based on the result of step 8 on page 7.

If the system-level diagnostics...	Then...
Were completed without any failures	<p>a. Clear the status logs: <code>sldiag device clearstatus</code></p> <p>b. Verify that the log is cleared: <code>sldiag device status</code> The following default response is displayed:</p> <pre style="background-color: #f0f0f0; padding: 5px;">SLDIAG: No log messages are present.</pre> <p>c. Exit Maintenance mode on both controllers: <code>halt</code> Attention: You must exit Maintenance mode on both controllers before proceeding any further.</p> <p>d. Boot up both controllers at the LOADER prompt: <code>boot_ontap</code></p> <p>e. If your system is a:</p> <ul style="list-style-type: none"> • 7-Mode system, enter <code>cf enable</code> • System running clustered Data ONTAP with two nodes in the cluster, enter <code>node::> cluster ha modify -configured true</code> <code>node::> storage failover modify -node node0 -enabled true</code> • System running clustered Data ONTAP with more than two nodes in the cluster, enter <code>node::> storage failover modify -node node0 -enabled true</code> <p>You have completed the system-level diagnostics.</p>
Resulted in some test failures	<p>Determine the cause of the problem:</p> <p>a. Exit Maintenance mode: <code>halt</code></p> <p>b. Perform a clean shutdown, and disconnect the power supplies.</p> <p>c. Verify that you have observed all the considerations identified for running the system-level diagnostics, that cables are securely connected, and that the hardware components are properly installed in the storage system.</p> <p>d. Reconnect the power supplies, and power on the storage system.</p> <p>e. Rerun the system-level diagnostics.</p>

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

Copyright © 1994–2017 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

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, 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. A current list of NetApp trademarks is available on the web.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto: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.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

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