



Chassis

ONTAP Systems

NetApp
June 24, 2022

Table of Contents

- Chassis 1
 - Overview of chassis replacement - AFF A700 and FAS9000 1
 - Shut down the controllers - AFF A700 and FAS9000 1
 - Move and replace hardware - AFF A700 and FAS9000 4
 - Complete the restoration and replacement process - AFF A700 and FAS9000 12

Chassis

Overview of chassis replacement - AFF A700 and FAS9000

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

- You can use this procedure with all versions of ONTAP supported by your system.
- This procedure is disruptive. For a two-node cluster, you will have a complete service outage and a partial outage in a multi-node cluster.

Shut down the controllers - AFF A700 and FAS9000

To replace the chassis, you must shutdown the controllers.

Option 1: Shut down the controllers

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

About this task

- If you have a cluster with more than two controllers, it must be in quorum. If the cluster is not in quorum or a healthy controller shows `false` for eligibility and health, you must correct the issue before shutting down the impaired controller; see the [Administration overview with the CLI](#).
- If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message  
MAINT=number_of_hours_downh
```

The following AutoSupport message suppresses automatic case creation for two hours: `cluster1:*>`

```
system node autosupport invoke -node * -type all -message MAINT=2h
```

Steps

1. If your system has two controller modules, disable the HA pair.

If your system is running clustered ONTAP with...	Then...
Two controllers in the cluster	<pre>cluster ha modify -configured false storage failover modify -node node0 -enabled false</pre>
More than two controllers in the cluster	<pre>storage failover modify -node node0 -enabled false</pre>

2. Halt the controller, pressing `y` when you are prompted to confirm the halt: `system node halt -node node_name`

The confirmation message looks like the following:

Warning: This operation will cause controller "node-name" to be marked as unhealthy. Unhealthy nodes do not participate in quorum voting. If the controller goes out of service and one more controller goes out of service there will be a data serving failure for the entire cluster. This will cause a client disruption. Use "cluster show" to verify cluster state. If possible bring other nodes online to improve the resiliency of this cluster.

Do you want to continue? {y|n}:



You must perform a clean system shutdown before replacing the chassis to avoid losing unwritten data in the nonvolatile memory (NVMEM/NVRAM). Depending on your system, if the NVMEM/NVRAM LED is flashing, there is content in the NVMEM/NVRAM that has not been saved to disk. You need to reboot the controller and start from the beginning of this procedure. If repeated attempts to cleanly shut down the controller fail, be aware that you might lose any data that was not saved to disk.

3. Where applicable, halt the second controller to avoid a possible quorum error message in an HA pair configuration: `system node halt -node second_node_name -ignore-quorum-warnings true -skip-lif-migration-before-shutdown true`

Answer *y* when prompted.

Option 2: Shut down a node in a two-node MetroCluster configuration

To shut down the impaired controller, you must determine the status of the controller and, if necessary, switch over the controller so that the healthy controller continues to serve data from the impaired controller storage.

About this task

- If you are using NetApp Storage Encryption, you must have reset the MSID using the instructions in the "Return a FIPS drive or SED to unprotected mode" section of [NetApp Encryption overview with the CLI](#).
- You must leave the power supplies turned on at the end of this procedure to provide power to the healthy controller.

Steps

1. Check the MetroCluster status to determine whether the impaired controller has automatically switched over to the healthy controller: `metrocluster show`
2. Depending on whether an automatic switchover has occurred, proceed according to the following table:

If the impaired controller...	Then...
Has automatically switched over	Proceed to the next step.
Has not automatically switched over	Perform a planned switchover operation from the healthy controller: <code>metrocluster switchover</code>

If the impaired controller...	Then...
Has not automatically switched over, you attempted switchover with the <code>metrocluster switchover</code> command, and the switchover was vetoed	Review the veto messages and, if possible, resolve the issue and try again. If you are unable to resolve the issue, contact technical support.

3. Resynchronize the data aggregates by running the `metrocluster heal -phase aggregates` command from the surviving cluster.

```
controller_A_1::> metrocluster heal -phase aggregates
[Job 130] Job succeeded: Heal Aggregates is successful.
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

4. Verify that the operation has been completed by using the `metrocluster operation show` command.

```
controller_A_1::> metrocluster operation show
  Operation: heal-aggregates
  State: successful
Start Time: 7/25/2016 18:45:55
End Time: 7/25/2016 18:45:56
Errors: -
```

5. Check the state of the aggregates by using the `storage aggregate show` command.

```
controller_A_1::> storage aggregate show
Aggregate      Size Available Used% State   #Vols  Nodes           RAID
Status
-----
...
aggr_b2       227.1GB   227.1GB   0% online    0  mcc1-a2
raid_dp, mirrored, normal...
```

6. Heal the root aggregates by using the `metrocluster heal -phase root-aggregates` command.

```
mcc1A::> metrocluster heal -phase root-aggregates
[Job 137] Job succeeded: Heal Root Aggregates is successful
```

If the healing is vetoed, you have the option of reissuing the `metrocluster heal` command with the `-override-vetoes` parameter. If you use this optional parameter, the system overrides any soft vetoes that prevent the healing operation.

7. Verify that the heal operation is complete by using the `metrocluster operation show` command on the destination cluster:

```
mccl1A::> metrocluster operation show
  Operation: heal-root-aggregates
    State: successful
  Start Time: 7/29/2016 20:54:41
    End Time: 7/29/2016 20:54:42
    Errors: -
```

8. On the impaired controller module, disconnect the power supplies.

Move and replace hardware - AFF A700 and FAS9000

Move the fans, hard drives, and controller module or modules from the impaired chassis to the new chassis, and swap out the impaired chassis from the equipment rack or system cabinet with the new chassis of the same model as the impaired chassis.

Step 1: Remove the power supplies


Steps

Removing the power supplies when replacing a chassis involves turning off, disconnecting, and then removing the power supply from the old chassis.

1. If you are not already grounded, properly ground yourself.
2. Turn off the power supply and disconnect the power cables:
 - a. Turn off the power switch on the power supply.
 - b. Open the power cable retainer, and then unplug the power cable from the power supply.
 - c. Unplug the power cable from the power source.
3. Press and hold the orange button on the power supply handle, and then pull the power supply out of the chassis.



When removing a power supply, always use two hands to support its weight.



1	Locking button
----------	----------------

4. Repeat the preceding steps for any remaining power supplies.

Step 2: Remove the fans


To remove the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

Steps

1. Remove the bezel (if necessary) with two hands, by grasping the openings on each side of the bezel, and then pulling it toward you until the bezel releases from the ball studs on the chassis frame.
2. Press the orange button on the fan module and pull the fan module straight out of the chassis, making sure that you support it with your free hand.



The fan modules are short. Always support the bottom of the fan module with your free hand so that it does not suddenly drop free from the chassis and injure you.



1	Orange release button
----------	-----------------------


3. Set the fan module aside.
4. Repeat the preceding steps for any remaining fan modules.

Step 3: Remove the controller module

To replace the chassis, you must remove the controller module or modules from the old chassis.

Steps

1. Unplug the cables from the impaired controller module, and keep track of where the cables were connected.
2. Slide the orange button on the cam handle downward until it unlocks.



<p>1</p>	<p>Cam handle release button</p>
<p>2</p>	<p>Cam handle</p>

3. Rotate the cam handle so that it completely disengages the controller module from the chassis, and then slide the controller module out of the chassis.

Make sure that you support the bottom of the controller module as you slide it out of the chassis.

4. Set the controller module aside in a safe place, and repeat these steps if you have another controller module in the chassis.

Step 4: Remove the I/O modules

Steps

To remove I/O modules from the old chassis, including the NVRAM modules, follow the specific sequence of steps. You do not have to remove the FlashCache module from the NVRAM module when moving it to a new

chassis.

1. Unplug any cabling associated with the target I/O module.

Make sure that you label the cables so that you know where they came from.

2. Remove the target I/O module from the chassis:
 - a. Depress the lettered and numbered cam button.


The cam button moves away from the chassis.

- b. Rotate the cam latch down until it is in a horizontal position.

The I/O module disengages from the chassis and moves about 1/2 inch out of the I/O slot.

- c. Remove the I/O module from the chassis by pulling on the pull tabs on the sides of the module face.

Make sure that you keep track of which slot the I/O module was in.



1	Lettered and numbered I/O cam latch
2	I/O cam latch completely unlocked


3. Set the I/O module aside.
4. Repeat the preceding step for the remaining I/O modules in the old chassis.

Step 5: Remove the De-stage Controller Power Module

Steps

You must remove the de-stage controller power modules from the old chassis in preparation for installing the replacement chassis.

1. Press the orange locking button on the module handle, and then slide the DCPM module out of the chassis.



1

DCPM module orange locking button

2. Set the DCPM module aside in a safe place and repeat this step for the remaining DCPM module.

Step 6: Replace a chassis from within the equipment rack or system cabinet

Steps

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

1. Remove the screws from the chassis mount points.



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 old chassis off the rack rails in a system cabinet or L brackets in an equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two or three people, install the replacement 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.

5. Slide the chassis all the way into the equipment rack or system cabinet.
6. Secure the front of the chassis to the equipment rack or system cabinet, using the screws you removed from the old chassis.
7. Secure the rear of the chassis to the equipment rack or system cabinet.
8. If you are using the cable management brackets, remove them from the old chassis, and then install them on the replacement chassis.
9. If you have not already done so, install the bezel.

Step 7: Move the USB LED module to the new chassis

Steps

Once the new chassis is installed into the rack or cabinet, you must move the USB LED module from the old chassis to the new chassis.

1. Locate the USB LED module on the front of the old chassis, directly under the power supply bays.
2. Press the black locking button on the right side of the module to release the module from the chassis, and then slide it out of the old chassis.
3. Align the edges of the module with the USB LED bay at the bottom-front of the replacement chassis, and gently push the module all the way into the chassis until it clicks into place.

Step 8: Install the de-stage controller power module when replacing the chassis

Steps

Once the replacement chassis is installed into the rack or system cabinet, you must reinstall the de-stage controller power modules into it.

1. Align the end of the DCPM module with the chassis opening, and then gently slide it into the chassis until it clicks into place.



The module and slot are keyed. Do not force the module into the opening. If the module does not go in easily, realign the module and slide it into the chassis.

2. Repeat this step for the remaining DCPM module.

Step 9: Install fans into the chassis

Steps

To install the fan modules when replacing the chassis, you must perform a specific sequence of tasks.

1. Align the edges of the replacement fan module with the opening in the chassis, and then slide it into the chassis until it snaps into place.

When inserted into a live system, the amber Attention LED flashes four times when the fan module is successfully inserted into the chassis.

2. Repeat these steps for the remaining fan modules.
3. Align the bezel with the ball studs, and then gently push the bezel onto the ball studs.

Step 10: Install I/O modules

Steps

To install I/O modules, including the NVRAM/FlashCache modules from the old chassis, follow the specific sequence of steps.

You must have the chassis installed so that you can install the I/O modules into the corresponding slots in the new chassis.

1. After the replacement chassis is installed in the rack or cabinet, install the I/O modules into their corresponding slots in the replacement chassis by gently sliding the I/O module into the slot until the lettered and numbered I/O cam latch begins to engage, and then push the I/O cam latch all the way up to lock the module in place.
2. Recable the I/O module, as needed.
3. Repeat the preceding step for the remaining I/O modules that you set aside.



If the old chassis has blank I/O panels, move them to the replacement chassis at this time.

Step 11: Install the power supplies

Steps

Installing the power supplies when replacing a chassis involves installing the power supplies into the replacement chassis, and connecting to the power source.

1. Using both hands, support and align the edges of the power supply with the opening in the system chassis, and then gently push the power supply into the chassis until it locks into place.

The power supplies are keyed and can only be installed one way.



Do not use excessive force when sliding the power supply into the system. You can damage the connector.

2. Reconnect the power cable and secure it to the power supply using the power cable locking mechanism.



Only connect the power cable to the power supply. Do not connect the power cable to a power source at this time.

3. Repeat the preceding steps for any remaining power supplies.

Step 12: Install the controller

Steps

After you install the controller module and any other components into the new chassis, boot it to a state where you can run the interconnect diagnostic test.

1. Align the end of the controller module with the opening in the chassis, and then gently push the controller module halfway into the system.



Do not completely insert the controller module in the chassis until instructed to do so.

2. Recable the console to the controller module, and then reconnect the management port.
3. Connect the power supplies to different power sources, and then turn them on.
4. With the cam handle in the open position, slide the controller module into the chassis and firmly push the controller module in until it meets the midplane and is fully seated, and then close the cam handle until it clicks into the locked position.



Do not use excessive force when sliding the controller module into the chassis; you might damage the connectors.

The controller module begins to boot as soon as it is fully seated in the chassis.

5. Repeat the preceding steps to install the second controller into the new chassis.
6. Boot each node to Maintenance mode:
 - a. As each node starts the booting, press `Ctrl-C` to interrupt the boot process when you see the message `Press Ctrl-C for Boot Menu`.



If you miss the prompt and the controller modules boot to ONTAP, enter `halt`, and then at the `LOADER` prompt enter `boot_ontap`, press `Ctrl-C` when prompted, and then repeat this step.

- b. From the boot menu, select the option for Maintenance mode.

Complete the restoration and replacement process - AFF A700 and FAS9000

You must verify the HA state of the chassis, run diagnostics, and return the failed part to NetApp, as described in the RMA instructions shipped with the kit.

Step 1: Verify and set 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.

Steps

1. In Maintenance mode, from either controller module, 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. If the displayed system state for the chassis 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`
- `mcc`
- `mcc-2n`

- `mccip`
- `non-ha`

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. Exit Maintenance mode: `halt`

The LOADER prompt appears.

Step 2: Running system-level diagnostics

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

Your system must be at the LOADER prompt to start System Level Diagnostics.

All commands in the diagnostic procedures are issued from the node where the component is being replaced.

Steps

1. If the node to be serviced is not at the LOADER prompt, perform the following steps:
 - a. Select the Maintenance mode option from the displayed menu.
 - b. After the node boots to Maintenance mode, halt the node: `halt`

After you issue the command, you should wait until the system stops at the LOADER prompt.



During the boot process, you can safely respond `y` to prompts:

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



Both controllers must be in Maintenance mode to run the interconnect test.

3. At the LOADER prompt, access the special drivers specifically designed for system-level diagnostics to function properly: `boot_diags`

During the boot process, you can safely respond `y` to the prompts until the Maintenance mode prompt (`*>`) appears.

4. Enable the interconnect diagnostics tests from the Maintenance mode prompt: `sldiag device modify -dev interconnect -sel enable`

The interconnect tests are disabled by default and must be enabled to run separately.

5. Run the interconnect diagnostics test from the Maintenance mode prompt: `sldiag device run -dev interconnect`


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

6. Verify that no hardware problems resulted from the replacement of the chassis: `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.

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

If the system-level diagnostics tests...	Then...
Were completed without any failures	<p>a. Clear the status logs: <code>sldiag device clearstatus</code></p> <p>b. Verify that the log was cleared: <code>sldiag device status</code></p> <p>The following default response is displayed:</p> <div data-bbox="672 512 1487 611" style="border: 1px solid #ccc; padding: 10px; background-color: #f9f9f9;"><pre>SLDIAG: No log messages are present.</pre></div> <p>c. Exit Maintenance mode on both controllers: <code>halt</code></p> <p>The system displays the LOADER prompt.</p> <div data-bbox="699 793 760 856" style="border: 1px solid #ccc; border-radius: 50%; width: 30px; height: 30px; display: flex; align-items: center; justify-content: center; margin: 10px 0;"></div> <p style="margin-left: 40px;">You must exit Maintenance mode on both controllers before proceeding any further.</p> <p>d. Enter the following command on both controllers at the LOADER prompt: <code>bye</code></p> <p>e. Return the node to normal operation:</p>
With two nodes in the cluster	<p>Issue these commands: <code>node::> cluster ha modify -configured true</code></p> <p><code>node::> storage failover modify -node node0 -enabled true</code></p>
With more than two nodes in the cluster	<p>Issue this command: <code>node::> storage failover modify -node node0 -enabled true</code></p>
In a two-node MetroCluster configuration	<p>Proceed to the next step.</p> <p>The MetroCluster switchback procedure is done in the next task in the replacement process.</p>
In a stand-alone configuration	<p>You have no further steps in this particular task.</p> <p>You have completed system-level diagnostics.</p>

If the system-level diagnostics tests...	Then...
Resulted in some test failures	Determine the cause of the problem. <ol style="list-style-type: none"> a. Exit Maintenance mode: <code>halt</code> b. Perform a clean shutdown, and then disconnect the power supplies. c. Verify that you have observed all of 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 then power on the storage system. e. Rerun the system-level diagnostics test.

Step 3: Switch back aggregates in a two-node MetroCluster configuration

After you have completed the FRU replacement in a two-node MetroCluster configuration, you can perform the MetroCluster switchback operation. This returns the configuration to its normal operating state, with the sync-source storage virtual machines (SVMs) on the formerly impaired site now active and serving data from the local disk pools.

This task only applies to two-node MetroCluster configurations.

Steps

1. Verify that all nodes are in the enabled state: `metrocluster node show`

```
cluster_B::> metrocluster node show

DR                               Configuration  DR
Group Cluster Node              State          Mirroring Mode
-----
1      cluster_A
      controller_A_1 configured    enabled    heal roots
completed
      cluster_B
      controller_B_1 configured    enabled    waiting for
switchback recovery
2 entries were displayed.
```

2. Verify that resynchronization is complete on all SVMs: `metrocluster vserver show`
3. Verify that any automatic LIF migrations being performed by the healing operations were completed successfully: `metrocluster check lif show`

4. Perform the switchback by using the `metrocluster switchback` command from any node in the surviving cluster.
5. Verify that the switchback operation has completed: `metrocluster show`

The switchback operation is still running when a cluster is in the `waiting-for-switchback` state:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured      switchover
Remote: cluster_A configured      waiting-for-switchback
```

The switchback operation is complete when the clusters are in the `normal` state.:

```
cluster_B::> metrocluster show
Cluster          Configuration State      Mode
-----
Local: cluster_B configured      normal
Remote: cluster_A configured      normal
```

If a switchback is taking a long time to finish, you can check on the status of in-progress baselines by using the `metrocluster config-replication resync-status show` command.

6. Reestablish any SnapMirror or SnapVault configurations.

Step 4: Return the failed part to NetApp

Return the failed part to NetApp, as described in the RMA instructions shipped with the kit. See the [Part Return & Replacements](#) page for further information.

Copyright Information

Copyright © 2022 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

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.