



Chassis

ONTAP Systems

NetApp
February 15, 2022

Table of Contents

- Chassis..... 1
 - Replace the chassis - AFF A400 1
 - Shut down the controllers - AFF A400 1
 - Move and replace hardware - AFF A400 4
 - Complete the restoration and replacement process - AFF A400 6

Chassis

Replace the chassis - AFF A400

To replace the chassis, you must move the fans and controller modules from the impaired chassis to the new chassis of the same model as the impaired chassis.

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 multinode cluster.

Shut down the controllers - AFF A400

Option 1: Shut down the controllers when replacing a chassis

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

About this task

- If you have a cluster with more than two nodes, it must be in quorum. If the cluster is not in quorum or a healthy node shows `false` for eligibility and health, you must correct the issue before shutting down the impaired node; 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 nodes in the cluster	<code>cluster ha modify -configured false storage failover modify -node node0 -enabled false</code>
More than two nodes in the cluster	<code>storage failover modify -node node0 -enabled false</code>

2. Halt the node, 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 node "node-name" to be marked as unhealthy. Unhealthy nodes do not participate in quorum voting. If the node goes out of service and one more node 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 node and start from the beginning of this procedure. If repeated attempts to cleanly shut down the node fail, be aware that you might lose any data that was not saved to disk.

3. Where applicable, halt the second node 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 node, you must determine the status of the node and, if necessary, switch over the node so that the healthy node continues to serve data from the impaired node 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 node.

Steps

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

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

If the impaired node...	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:

```
mcclA::> 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 A400

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 controller modules

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

1. If you are not already grounded, properly ground yourself.
2. Release the power cable retainers, and then unplug the cables from the power supplies.
3. Loosen the hook and loop strap binding the cables to the cable management device, and then unplug the system cables and SFPs (if needed) from the controller module, keeping track of where the cables were connected.

Leave the cables in the cable management device so that when you reinstall the cable management device, the cables are organized.

4. Remove and set aside the cable management devices from the left and right sides of the controller module.
5. Press down on both of the locking latches, and then rotate both latches downward at the same time.

The controller module moves slightly out of the chassis.

6. 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.

7. Set the controller module aside in a safe place, and repeat these steps for the other controller module in the chassis.

Step 2: Move the fans

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

1. If you are not already grounded, properly ground yourself.
2. 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.
3. Press down the release latch on the fan module cam handle, and then rotate the cam handle downward.

The fan module moves a little bit away from the chassis.

4. Pull the fan module straight out from the chassis, making sure that you support it with your free hand so that it does not swing out of the chassis.



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.

5. Set the fan module aside.
6. Repeat the preceding steps for any remaining fan modules.
7. Insert the fan module into the replacement chassis by aligning it with the opening, and then sliding it into the chassis.
8. Push firmly on the fan module cam handle so that it is seated all the way into the chassis.

The cam handle raises slightly when the fan module is completely seated.

9. Swing the cam handle up to its closed position, making sure that the cam handle release latch clicks into the locked position.

The fan LED should be green after the fan is seated and has spun up to operational speed.

10. Repeat these steps for the remaining fan modules.

Step 3: Replace a chassis from within the equipment rack or system cabinet

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.
2. With two people, slide the old chassis off the rack rails in a system cabinet or equipment rack, and then set it aside.
3. If you are not already grounded, properly ground yourself.
4. Using two 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 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. If you have not already done so, install the bezel.

Step 4: Install the controller modules

After you install the controller modules into the new chassis, you need to boot it to a state where you can run the diagnostic test.

For HA pairs with two controller modules in the same chassis, the sequence in which you install the controller module is especially important because it attempts to reboot as soon as you completely seat it in the chassis.

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. Complete the installation of the controller module:
 - a. Plug the power cord into the power supply, reinstall the power cable locking collar, and then connect the power supply to the power source.
 - b. Using the locking latches, firmly push the controller module into the chassis until the locking latches begin to rise.



Do not use excessive force when sliding the controller module into the chassis to avoid damaging the connectors.

- c. Fully seat the controller module in the chassis by rotating the locking latches upward, tilting them so that they clear the locking pins, gently push the controller all the way in, and then lower the locking latches into the locked position.

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

- d. If you have not already done so, reinstall the cable management device.
- e. Interrupt the normal boot process and boot to LOADER by pressing `Ctrl-C`.



If your system stops at the boot menu, select the option to boot to LOADER.

- f. At the LOADER prompt, enter `bye` to reinitialize the PCIe cards and other components.
- g. Interrupt the boot process and boot to the LOADER prompt by pressing `Ctrl-C`.

If your system stops at the boot menu, select the option to boot to LOADER.

4. Repeat the preceding steps to install the second controller into the new chassis.

Complete the restoration and replacement process - AFF A400

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.

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. Reinstall the bezel on the front of the system.

Step 2: Run diagnostics

After you have replaced a component in your system, you should run diagnostic tests on that component.

Your system must be at the LOADER prompt to start diagnostics.

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

1. If the node to be serviced is not at the LOADER prompt, reboot the node: `system node halt -node node_name`

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

2. At the LOADER prompt, access the special drivers specifically designed for system-level diagnostics to function properly: `boot_diags`
3. Select **Scan System** from the displayed menu to enable running the diagnostics tests.
4. Select **Test system** from the displayed menu to run diagnostics tests.
5. Select the test or series of tests from the various sub-menus.
6. Proceed based on the result of the preceding step:
 - If the test failed, correct the failure, and then rerun the test.
 - If the test reported no failures, select Reboot from the menu to reboot the system.

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

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 © 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.