NetApp

Chassis

Install and maintain

NetApp
February 13, 2026

Table of Contents

Chassis

Overview of chassis replacement - AFF C400

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 C400

Shut down or take over the impaired controller using the appropriate procedure for your configuration.

Option 1: Shut down the controllers when replacing a chassis

This procedure is for systems with two node configurations. For more information about graceful shutdown when servicing a cluster, see Gracefully shutdown and power up your storage system Resolution Guide - NetApp Knowledge Base.

Before you begin

- Make sure you have the necessary permissions and credentials:
    - Local administrator credentials for ONTAP.
    - BMC accessability for each controller.
- Make sure you have the necessary tools and equipment for the replacement.
- As a best practice before shutdown, you should:
    - Perform additional system health checks.
    - Upgrade ONTAP to a recommended release for the system.
    - Resolve any Active IQ Wellness Alerts and Risks. Make note of any faults presently on the system, such as LEDs on the system components.

Steps

1. Log into the cluster through SSH or log in from any node in the cluster using a local console cable and a laptop/console.
2. Stop all clients/host from accessing data on the NetApp system.
3. Suspend external backup jobs.
4. If AutoSupport is enabled, suppress case creation and indicate how long you expect the system to be offline:

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

5. Identify the SP/BMC address of all cluster nodes:

```
system service-processor show -node * -fields address
```

6. Exit the cluster shell:

```
exit
```

7. Log into SP/BMC over SSH using the IP address of any of the nodes listed in the output from the previous step to monitor progress.

   If you are using a console/laptop, log into the controller using the same cluster administrator credentials.

8. Halt the two nodes located in the impaired chassis:

```
system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown
true -ignore-quorum-warnings true -inhibit-takeover true
```

> (i) For clusters using SnapMirror synchronous operating in StrictSync mode: `system node halt -node <node1>,<node2> -skip-lif-migration-before-shutdown true -ignore-quorum-warnings true -inhibit-takeover true -ignore-strict -sync-warnings true`

9. Enter **y** for each controller in the cluster when you see:

```
Warning: Are you sure you want to halt node <node_name>? {y|n}:
```

10. Wait for each controller to halt and display the LOADER prompt.

Option 2: Shut down a controller 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

- 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: `metrocluster switchover`

If the impaired controller…	Then…
Has not automatically switched over, you attempted switchover with the `metrocluster switchover` 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:

```
mcc1A::> 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.

Replace hardware - AFF C400

Move the fans, hard drives, and controller module from the impaired chassis to the new chassis, and swap out the impaired chassis 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.

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.

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.

   (i) 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.

      (i) 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`.

      (i) 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 C400

You must verify the HA state of the chassis 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: Switch back aggregates in a two-node MetroCluster configuration

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 3: 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 and Replacements page for further information.