



Controller

Install and maintain

NetApp
July 01, 2024

Table of Contents

- Controller 1
 - Overview of controller module replacement - AFF A1K 1
 - Shut down the impaired controller - AFF A1K 1
 - Replace the controller module hardware - AFF A1K 3
 - Restore and verify the system configuration - AFF A1K 8
 - Give back the controller - AFF A1K 9
 - Complete system restoration - AFF A1K 11

Controller

Overview of controller module replacement - AFF A1K

You must review the prerequisites for the replacement procedure and select the correct one for your version of the ONTAP operating system.

- All drive shelves must be working properly.
- The healthy controller must be able to take over the controller that is being replaced (referred to in this procedure as the “impaired controller”).
- If your system is in a MetroCluster configuration, you must review the section [Choosing the correct recovery procedure](#) to determine whether you should use this procedure.
- You must replace the failed component with a replacement FRU component you received from your provider.
- You must be replacing a controller module with a controller module of the same model type. You cannot upgrade your system by just replacing the controller module.
- You cannot change any drives or drive shelves as part of this procedure.
- Because the boot device is located on the System Management module that is installed in the back of the system, you do not need to move the boot device when replacing a controller module.
- It is important that you apply the commands in these steps on the correct systems:
 - The *impaired* controller is the controller that is being replaced.
 - The *replacement* controller is the new controller that is replacing the impaired controller.
 - The *healthy* controller is the surviving controller.
- You must always capture the controller’s console output to a text log file.

This provides you a record of the procedure so that you can troubleshoot any issues that you might encounter during the replacement process.

Shut down the impaired controller - AFF A1K

Option 1: Most systems

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

About this task

- If you have a SAN system, you must have checked event messages (`cluster kernel-service show`) for the impaired controller SCSI blade. The `cluster kernel-service show` command (from `priv advanced` mode) displays the node name, quorum status of that node, availability status of that node, and operational status of that node.

Each SCSI-blade process should be in quorum with the other nodes in the cluster. Any issues must be resolved before you proceed with the replacement.

- 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 controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).

Steps

1. If AutoSupport is enabled, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=<# of hours>h
```

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

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

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`



When you see *Do you want to disable auto-giveback?*, enter *y*.

3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <i>y</i> when prompted.
System prompt or password prompt	Take over or halt the impaired controller from the healthy controller: <code>storage failover takeover -ofnode <i>impaired_node_name</i></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <i>y</i> .

Option 2: Controller is in a MetroCluster

To shut down the impaired controller, you must determine the status of the controller and, if necessary,

take over the controller so that the healthy controller continues to serve data from the impaired controller storage.

- 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 controller shows false for eligibility and health, you must correct the issue before shutting down the impaired controller; see [Synchronize a node with the cluster](#).
- If you have a MetroCluster configuration, you must have confirmed that the MetroCluster Configuration State is configured and that the nodes are in an enabled and normal state (`metrocluster node show`).

Steps

1. 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
```

2. Disable automatic giveback from the console of the healthy controller: `storage failover modify -node local -auto-giveback false`
3. Take the impaired controller to the LOADER prompt:

If the impaired controller is displaying...	Then...
The LOADER prompt	Go to the next step.
Waiting for giveback...	Press Ctrl-C, and then respond <code>y</code> when prompted.
System prompt or password prompt (enter system password)	Take over or halt the impaired controller from the healthy controller: <code>storage failover takeover -ofnode <i>impaired_node_name</i></code> When the impaired controller shows Waiting for giveback..., press Ctrl-C, and then respond <code>y</code> .


Replace the controller module hardware - AFF A1K

To replace the controller, you must remove the impaired controller, move FRU components to the replacement controller module, install the replacement controller module in the enclosure, and then boot the system to Maintenance mode.

Step 1 Remove the controller module

You must remove the controller module from the enclosure when you replace the controller module or replace a component inside the controller module.

1. Check the NVRAM status LED located in slot 4/5 of the system. There is also an NVRAM LED on the front panel of the controller module. Look for the NV icon:



1	NVRAM status LED
2	NVRAM attention LED




If the NVRAM status LED is flashing, it could mean the controller module was not taken over or halted properly (uncommitted data). If the impaired controller module was not successfully taken over by the partner controller module, contact [NetApp Support](#) before continuing with this procedure.

The general behavior of the NVRAM status LED on the impaired controller module is as follows:

- The NVRAM status LED flashes when power is removed from the controller module and the controller module is in the "waiting for giveback" state, or the controller module is not taken over or halted properly (uncommitted data).
 - The NVRAM status LED flashes when the controller module is removed from the enclosure and could mean the controller module is not taken over or halted properly (uncommitted data). Confirm that the controller module has been cleanly takeover by the partner controller module or the impaired controller module shows `waiting for giveback`. Then, the flashing LED can be ignored (and the controller can be removed from the enclosure).
2. If you are not already grounded, properly ground yourself.
 3. On the front of the unit, hook your fingers into the holes in the locking cams, squeeze the tabs on the cam levers, and gently, but firmly rotate both latches toward you at the same time.

The controller module moves slightly out of the enclosure.



1	a Locking cam latches
----------	-----------------------

4. Slide the controller module out of the enclosure and place it on a flat, stable surface.

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


Step 2: Move the fans

You must remove the five fan modules from the impaired controller module to the replacement controller module.

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 the gray locking 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	Black locking button
----------	----------------------

4. Install the fan in the replacement controller module:


- a. Align the edges of the fan housing with the opening in the front of the replacement controller module.
- b. Gently slide the fan module all the way into the replacement controller module until it locks in place.

5. Repeat the preceding steps for the remaining fan modules.

Step 3: Move the NV battery

Move the NV battery to the replacement controller.

1. Open the NV battery air duct cover and locate the NV battery.




1	NV battery air duct cover
2	NV battery plug
3	NV battery pack

2. Lift the battery up to access the battery plug.
3. Squeeze the clip on the face of the battery plug to release the plug from the socket, and then unplug the battery cable from the socket.
4. Lift the battery out of the air duct and controller module.
5. Move the battery pack to the replacement controller module and then install it in the NV battery air duct:
 - a. Open the NV battery air duct in the replacement controller module.
 - b. Plug the battery plug into the socket and make sure that the plug locks into place.
 - c. Insert the battery pack into the slot and press firmly down on the battery pack to make sure that it is locked into place.
 - d. Close the air duct cover.

Step 4: Move system DIMMs

Move the DIMMs to the replacement controller module.

1. Open the motherboard air duct and locate the DIMMs.



1	System DIMM
----------	-------------

2. Note the orientation of the DIMM in the socket so that you can insert the DIMM in the replacement controller module in the proper orientation.
3. Eject the DIMM from its slot by slowly pushing apart the two DIMM ejector tabs on either side of the DIMM, and then slide the DIMM out of the slot.



Carefully hold the DIMM by the edges to avoid pressure on the components on the DIMM circuit board.

4. Locate the slot where you are installing the DIMM in the replacement controller module.
5. Insert the DIMM squarely into the slot.

The DIMM fits tightly in the slot, but should go in easily. If not, realign the DIMM with the slot and reinsert it.



Visually inspect the DIMM to verify that it is evenly aligned and fully inserted into the slot.

6. Push carefully, but firmly, on the top edge of the DIMM until the ejector tabs snap into place over the notches at the ends of the DIMM.
7. Repeat these steps for the remaining DIMMs. Close the motherboard air duct.

Step 5: Install the controller module

Reinstall the controller module and boot it.

1. If you have not already done so, close the air duct.
2. Align the end of the controller module with the opening in the enclosure, and slide the controller module into the chassis with the levers rotated away from the front of the system.

3. Once the controller module stops you from sliding it farther, rotate the cam handles inward until they latch back under the fans



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

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

4. Restore automatic giveback if you disabled it by using the `storage failover modify -node local -auto-giveback true` command.
5. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Restore and verify the system configuration - AFF A1K

After completing the hardware replacement, you verify the low-level system configuration of the replacement controller and reconfigure system settings as necessary.

Step 1: Set and verify system time

You should check the time and date on the replacement controller module against the healthy controller module in an HA pair, or against a reliable time server in a stand-alone configuration. If the time and date do not match, you must reset them on the replacement controller module to prevent possible outages on clients due to time differences.

About this task

It is important that you apply the commands in the steps on the correct systems:

- The *replacement* node is the new node that replaced the impaired node as part of this procedure.
- The *healthy* node is the HA partner of the *replacement* node.

Steps

1. If the *replacement* node is not at the LOADER prompt, halt the system to the LOADER prompt.
2. On the *healthy* node, check the system time: `cluster date show`

The date and time are based on the configured timezone.

3. At the LOADER prompt, check the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

4. If necessary, set the date in GMT on the replacement node: `set date mm/dd/yyyy`
5. If necessary, set the time in GMT on the replacement node: `set time hh:mm:ss`
6. At the LOADER prompt, confirm the date and time on the *replacement* node: `show date`

The date and time are given in GMT.

Step 2: Verify and set the HA state of the chassis

You must verify the HA state of the controller module and, if necessary, update the state to match your system configuration.

1. Boot to Maintenance mode on the replacement controller module and verify that all components display the same HA state: `ha-config show`

The HA state should be the same for all components.

2. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`

The value for HA-state can be one of the following:

- `ha`
- `mcc` (not supported)
- `mccip`
- `non-ha` (not supported)

3. If the displayed system state of the controller module does not match your system configuration, set the HA state for the controller module: `ha-config modify controller ha-state`
4. Confirm that the setting has changed: `ha-config show`

Give back the controller - AFF A1K

Continue the replacement procedure by giving back the controller.

Step 1: Give back the controller

1. If your storage system has Encryption configured, you must restore Storage or Volume Encryption functionality using the following procedure to reboot the system:
 - a. Boot to Menu and run Option 10
 - b. Input the passphrase & backup up data, then do Normal boot see [Restore onboard key management encryption keys](#).
 - c. Perform CFO only giveback
 - d. Perform Onboard Sync and verify SVM-KEK is set to true see [Giveback after MB replacement fails - operation was vetoed by keymanager](#)
 - e. Giveback SFO, (no force)
2. If your system does not have Encryption configured, complete the following procedure to reboot the system:
 - a. Boot to Menu and run Option 1.
 - b. Give back the controller:
 - c. From the healthy controller, give back the replaced controller's storage: `storage failover giveback -ofnode replacement_node_name`

The *replacement* controller takes back its storage and completes booting.



If the giveback is vetoed, you can consider overriding the vetoes.

Find the [High-Availability Configuration content for your version of ONTAP 9](#)

d. After the giveback has been completed, confirm that the HA pair is healthy and that takeover is possible: `storage failover show`

3. If the system is in a MetroCluster configuration, monitor the status of the controller: `metrocluster node show`

The MetroCluster configuration takes a few minutes after the replacement to return to a normal state, at which time each controller will show a configured state, with DR Mirroring enabled and a mode of normal. The `metrocluster node show -fields node-systemid` command output displays the old system ID until the MetroCluster configuration returns to a normal state.

4. If the controller is in a MetroCluster configuration, depending on the MetroCluster state, verify that the DR home ID field shows the original owner of the disk if the original owner is a controller on the disaster site.

This is required if both of the following are true:

- The MetroCluster configuration is in a switchover state.
- The *replacement* controller is the current owner of the disks on the disaster site.

[Disk ownership changes during HA takeover and MetroCluster switchover in a four-node MetroCluster configuration](#)

5. If your system is in a MetroCluster configuration, verify that each controller is configured: `metrocluster node show - fields configuration-state`

```
node1_siteA::> metrocluster node show -fields configuration-state

dr-group-id          cluster node          configuration-state
-----
-----
1 node1_siteA        node1mcc-001         configured
1 node1_siteA        node1mcc-002         configured
1 node1_siteB        node1mcc-003         configured
1 node1_siteB        node1mcc-004         configured

4 entries were displayed.
```

6. Verify that the expected volumes are present for each controller: `vol show -node node-name`

7. If you disabled automatic takeover on reboot, enable it from the healthy controller: `storage failover modify -node replacement-node-name -onreboot true`

8. If AutoSupport is enabled, restore/unsuppress automatic case creation by using the `system node autosupport invoke -node * -type all -message MAINT=END` command.

Complete system restoration - AFF A1K

To restore your system to full operation, you must verify the LIFs, check cluster health, and return the failed part to NetApp.

Step 1: Verify LIFs and and check cluster health

Before returning the *replacement* node to service, you should verify that the LIFs are on their home ports, check the cluster health, and reset automatic giveback.

Steps

1. Verify that the logical interfaces are reporting to their home server and ports: `network interface show -is-home false`

If any LIFs are listed as false, revert them to their home ports: `network interface revert -vserver * -lif *`

2. Check the health of your cluster. See the [How to perform a cluster health check with a script in ONTAP KB](#) article for more information.
3. If automatic giveback was disabled, reenable it: `storage failover modify -node local -auto-giveback true`

Step 2: 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 © 2024 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

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.