



Disruptively transitioning when existing shelves are not supported on new controllers (ONTAP 9.8 and later)

ONTAP MetroCluster

Martin Houser, Zachary Wambold
April 26, 2021

Table of Contents

- Disruptively transitioning when existing shelves are not supported on new controllers (ONTAP 9.8 and later) 1
 - Requirements for transition when shelves are not supported on the new nodes 1
 - Workflow for disruptive transition when shelves are not supported by new controllers 2
 - Preparing the new controller modules 3
 - Attaching the new disk shelves to the existing MetroCluster FC controllers 3
 - Migrate root aggregates and move data to the new disk shelves 4
 - Transitioning the configuration 7

Disruptively transitioning when existing shelves are not supported on new controllers (ONTAP 9.8 and later)

Starting with ONTAP 9.8, you can disruptively transition a two-node MetroCluster FC configuration and move data from the existing drive shelves even if the existing storage shelves are not supported by the new MetroCluster IP nodes.

- This procedure should only be used if the existing storage shelf models are not supported by the new MetroCluster IP platform models.
- This procedure is supported on systems running ONTAP 9.8 and later.
- This procedure is disruptive.
- This procedure applies only to a two-node MetroCluster FC configuration.

If you have a four-node MetroCluster FC configuration, see [Choosing your transition procedure](#).

- You must meet all requirements and follow all steps in the procedure.

Requirements for transition when shelves are not supported on the new nodes

Before starting the transition process, you must make sure the configuration meets the requirements.

- The existing configuration must be a two-node fabric-attached or stretch MetroCluster configuration and all nodes must be running ONTAP 9.8 or later.

The new MetroCluster IP controller modules should be running the same version of ONTAP 9.8.

- The existing and new platforms must be a supported combination for transition.

[Supported platforms for nondisruptive transition](#)

- It must meet all requirements and cabling as described in the *MetroCluster Installation and Configuration Guides*.

[Fabric-attached MetroCluster installation and configuration](#)

- New storage shelves provided with the new controllers (node_A_1-IP, node_A_2-IP, node_B_1-IP and node_B_2-IP) must be supported by the old controllers (node_A_1-FC and node_B_1-FC).

[NetApp Hardware Universe](#)

- The old storage shelves are **not** supported by the new MetroCluster IP platform models.

[NetApp Hardware Universe](#)

- Depending on the spare disks available in the existing shelves, additional drives must be added.

This might require additional drive shelves.

You need to have additional 14 to 18 drives for each controller:

- Three pool0 drives
 - Three pool1 drives
 - Two spare drives
 - Six to ten drives for the system volume
- You must ensure that the configuration, including the new nodes, does not exceed the platform limits for the configuration, including drive count, root aggregate size capacity, etc.

This information is available for each platform model at *NetApp Hardware Universe*.

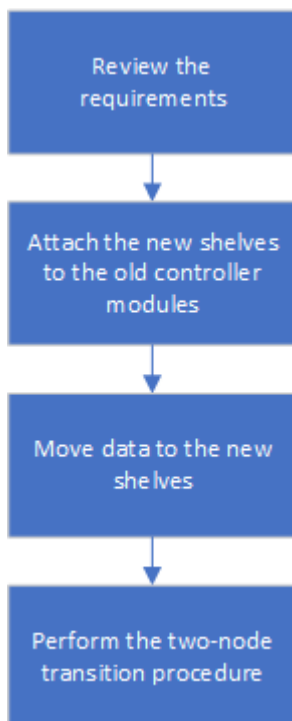
[NetApp Hardware Universe](#)

- You must have remote console access for all six nodes from either MetroCluster site or plan for travel between the sites as required by the procedure.

Workflow for disruptive transition when shelves are not supported by new controllers

If the existing shelf models are not supported by the new platform models, you must attach the new shelves to the old configuration, move data onto the new shelves, and then transition to the new configuration.

As you prepare for the transition, plan for travel between the sites. Note that after the remote nodes are racked and cabled, you need serial terminal access to the nodes. Service Processor access is not be available until the nodes are configured.



Preparing the new controller modules

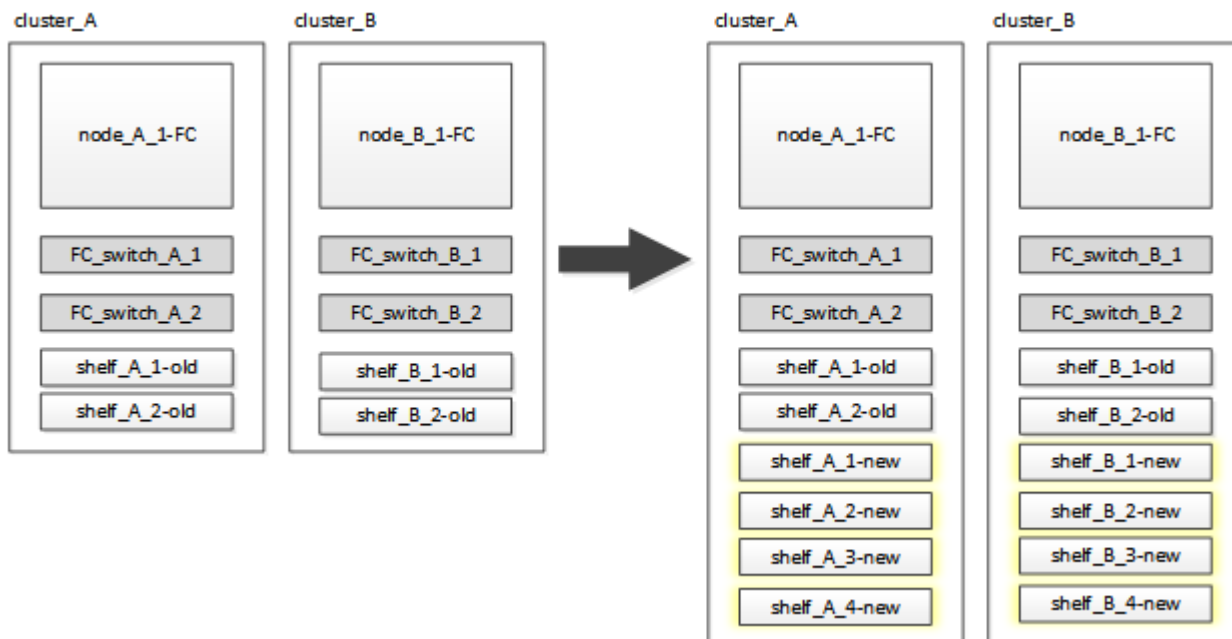
You must clear the configuration and disk ownership on the new controller modules and the new storage shelves.

1. With the new storage shelves attached to the new MetroCluster IP controller modules, perform all the steps in the section [Preparing the MetroCluster IP controllers](#).
2. Disconnect the new storage shelves from the new MetroCluster IP controller modules.

Attaching the new disk shelves to the existing MetroCluster FC controllers

You must attach the new drive shelves to the existing controller modules before transitioning to a MetroCluster IP configuration.

The following illustration shows the new shelves attached to the MetroCluster FC configuration.



1. Disable disk autoassignment on node_A_1-FC and node_A_2-FC: `disk option modify -node node-name -autoassign off`

This command must be issued on each node.

Disk auto-assign is disabled to avoid assignment of the shelves to be added to node_A_1-FC and node_B_1-FC. As part the transition, disks are needed for nodes node_A_1-IP and node_B_2-IP and if autoassign is allowed, disk ownership would later need to be removed before disks could be assigned to node_A_1-IP and node_B_2-IP.

2. Attach the new shelves to the existing MetroCluster FC nodes, using FC-to-SAS bridges, if necessary.

See the requirements and procedures in the described in the *MetroCluster Maintenance Guide*.

[Hot-adding storage to a MetroCluster FC configuration](#)

Migrate root aggregates and move data to the new disk shelves

You must move the root aggregates from the old drive shelves to the new drive shelves that will be used by the MetroCluster IP nodes.

This task is performed prior to the transition on the existing nodes (node_A_1-FC and node_B_1-FC).

1. Perform a negotiated switchover from controller node_B_1-FC: `metrocluster switchover`
2. Perform the heal aggregates and heal root steps of the recovery from node_B_1-FC: `metrocluster heal -phase aggregates` `metrocluster heal -phase root-aggregates`
3. Boot controller node_A_1-FC: `boot_ontap`
4. Assign the unowned disks on the new shelves to the appropriate pools for controller node_A_1-FC:
 - a. Identify the disks on the shelves: `disk show -shelf pool_0_shelf -fields container-type,diskpathnames` `disk show -shelf pool_1_shelf -fields container-type,diskpathnames`
 - b. Enter local mode so the commands are run on the local node: `run local`
 - c. Assign the disks: `disk assign disk1disk2disk3disk... -p 0` `disk assign disk4disk5disk6disk... -p 1`
 - d. Exit local mode: `exit`
5. Create a new mirrored aggregate to become the new root aggregate for controller node_A_1-FC:
 - a. Set the privilege mode to advanced: `set priv advanced`
 - b. Create the aggregate: `aggregate create -aggregate new_aggr -disklist disk1, disk2, disk3,... -mirror-disklist disk4disk5, disk6,... -raidtypesame-as -existing-root -force-small-aggregate true aggr show -aggregate new_aggr -fields percent-snapshot-space`

If the percent-snapshot-space value is less than 5 percent, you must increase it to a value higher than 5 percent: `aggr modify new_aggr -percent-snapshot-space 5`
 - c. Set the privilege mode back to admin: `set priv admin`
6. Confirm that the new aggregate is created properly: `node run -node local sysconfig -r`
7. Create the node and cluster-level configuration backups:



When the backups are created during switchover, the cluster is aware of the switched over state on recovery. You must ensure that the backup and upload of the system configuration is successful as without this backup it is **not** possible to reform the MetroCluster configuration between clusters.

- a. Create the cluster backup: `system configuration backup create -node local -backup -type cluster -backup-name cluster-backup-name`
- b. Check cluster backup creation `job show -id job-idstatus`
- c. Create the node backup: `system configuration backup create -node local -backup -type node -backup-name node-backup-name`

d. Check for both cluster and node backups: `system configuration backup show`

You can repeat the command until both backups are shown in the output.

8. Make copies of the backups.

The backups must be stored at a separate location because they will be lost locally when the new root volume is booted.

You can upload the backups to an FTP or HTTP server, or copy the backups using `scp` commands.

Method	
Upload the backup to the FTP or HTTP server	<p>a. Upload the cluster backup: <code>system configuration backup upload -node local -backup cluster-backup-name -destination URL</code></p> <p>b. Upload the node backup: <code>system configuration backup upload -node local -backup node-backup-name -destination URL</code></p>
Copy the backups onto a remote server using secure copy	<p>From the remote server use the following <code>scp</code> commands:</p> <p>a. Copy the cluster backup: <code>scp diag@node-mgmt-FC:/mroot/etc/backups/config/cluster-backup-name.7z .</code></p> <p>b. Copy the node backup: <code>scp diag@node-mgmt-FC:/mroot/etc/backups/config/node-backup-name.7z .</code></p>

9. Halt node_A_1-FC: `halt -node local -ignore-quorum-warnings true`

10. Boot node_A_1-FC to Maintenance mode: `boot_ontap maint`

11. From Maintenance mode, make required changes to set the aggregate as root:

a. Set the HA policy to cfo: `aggr options new_aggr ha_policy cfo`

Respond `yes` when prompted to proceed.

Are you sure you want to proceed (y/n)?

b. Set the new aggregate as root: `aggr options new_aggr root`

c. Halt to the LOADER prompt: `halt`

12. Boot the controller and back up the system configuration.

The node boots in recovery mode when the new root volume is detected

- a. Boot the controller: `boot_ontap`
- b. Log in and back up the configuration.

When you log in, you will see the following warning:

```
Warning: The correct cluster system configuration backup must be
restored. If a backup
from another cluster or another system state is used then the root
volume will need to be
recreated and NGS engaged for recovery assistance.
```

- c. Enter advanced privilege mode: `set -privilege advanced`
- d. Back up the cluster configuration to a server: `system configuration backup download -node local -source URL of server/cluster-backup-name.7z`
- e. Back up the node configuration to a server: `system configuration backup download -node local -source URL of server/node-backup-name.7z`
- f. Return to admin mode: `set -privilege admin`

13. Check the health of the cluster:

- a. Issue the following command: `cluster show`
- b. Set the privilege mode to advanced: `set -privilege advanced`
- c. Verify the cluster configuration details: `cluster ring show`
- d. Return to the admin privilege level: `set -privilege admin`

14. Confirm the operational mode of the MetroCluster configuration and perform a MetroCluster check.

- a. Confirm the MetroCluster configuration and that the operational mode is normal: `metrocluster show`
- b. Confirm that all expected nodes are shown: `metrocluster node show`
- c. Issue the following command: `metrocluster check run`
- d. Display the results of the MetroCluster check: `metrocluster check show`

15. Perform a switchback from controller node_B_1-FC: `metrocluster switchback`

16. Verify the operation of the MetroCluster configuration:

- a. Confirm the MetroCluster configuration and that the operational mode is normal: `metrocluster show`
- b. Perform a MetroCluster check: `metrocluster check run`
- c. Display the results of the MetroCluster check: `metrocluster check show`

17. Add the new root volume to the Volume Location Database.

- a. Set the privilege mode to advanced: `set -privilege advanced`
 - b. Add the volume to the node: `volume add-other-volumes -node node_A_1-FC`
 - c. Return to the admin privilege level: `set -privilege admin`
18. Check that the volume is now visible and has mroot.
- a. Display the aggregates: `storage aggregate show`
 - b. Verify that the root volume has mroot: `storage aggregate show -fields has-mroot`
 - c. Display the volumes: `volume show`
19. Create a new security certificate to re-enable access to System Manager: `security certificate create -common-name name -type server -size 2048`
20. Repeat the previous steps to migrate the aggregates on shelves owned by node_A_1-FC.
21. Perform a cleanup.

You must perform the following steps on both node_A_1-FC and node_B_1-FC to remove the old root volume and root aggregate.

- a. Delete the old root volume: `run localvol offline old_vol0vol destroy old_vol0 exitvolume remove-other-volume -vserver node_name -volume old_vol0`
 - b. Delete the original root aggregate: `aggr offline -aggregate old_aggr0_site``aggr delete -aggregate old_aggr0_site`
22. Migrate the data volumes to aggregates on the new controllers, one volume at a time.
- Use the following section of the *Controller Upgrade Express Guide*.
- [Creating an aggregate and moving volumes to the new nodes](#)
23. Retire the old shelves by performing all the steps in the section [Retiring shelves moved from node_A_1-FC and node_A_2-FC](#).

Transitioning the configuration

You must follow the detailed transition procedure.

In the following steps you are directed to other sections in this guide. You must perform the steps in each referenced section in the order given.

1. Plan port mapping.

Perform all the steps in the section [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#).

2. Prepare the MetroCluster IP controllers.

Perform all the steps in the section [Preparing the MetroCluster IP controllers](#).

3. Verify the health of the MetroCluster configuration.

Perform all the steps in the section [Verifying the health of the MetroCluster FC configuration](#)

4. Prepare and remove the existing MetroCluster FC nodes.

Perform all the steps in the section [Transitioning the MetroCluster FC nodes](#).

5. Add the new MetroCluster IP nodes.

Perform all the steps in the section [Connecting the MetroCluster IP controller modules](#).

6. Complete the transition and initial configuration of the new MetroCluster IP nodes.

Perform all the steps in the section [Configuring the new nodes and completing transition](#).

Copyright Information

Copyright © 2021 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.