



Disruptively transitioning from MetroCluster FC to MetroCluster IP when retiring storage shelves (ONTAP 9.8 and later)

ONTAP MetroCluster

Martin Houser, Zachary Wambold
April 19, 2021

Table of Contents

- Disruptively transitioning from MetroCluster FC to MetroCluster IP when retiring storage shelves (ONTAP 9.8 and later) 1
 - Requirements for transition when retiring old shelves 1
 - Workflow for disruptive transition when moving data and retiring old storage shelves 2
 - Transitioning the configuration 2
 - Migrating the root aggregates 3
 - Migrating the data aggregates 4
 - Retiring shelves moved from node_A_1-FC and node_A_2-FC 4
 - Completing transition 6

Disruptively transitioning from MetroCluster FC to MetroCluster IP when retiring storage shelves (ONTAP 9.8 and later)

Starting with ONTAP 9.8, you can disruptively transition a two-node MetroCluster FC configuration to a four-node MetroCluster IP configuration and retire the existing storage shelves. The procedure includes steps to move data from the existing drive shelves to the new configuration, and then retire the old shelves.

- This procedure is used when you plan to retire the existing storage shelves and move all data to the new shelves in the MetroCluster IP configuration.
- The existing storage shelf models must be supported by the new MetroCluster IP nodes.
- 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 retiring old shelves

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

- It 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](#)

The new configuration must also meet the following requirements:

- The new MetroCluster IP platform models must support the old storage shelf 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 - 18 drives for each controller:

- Three pool 0 drives
 - Three pool 1 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*.

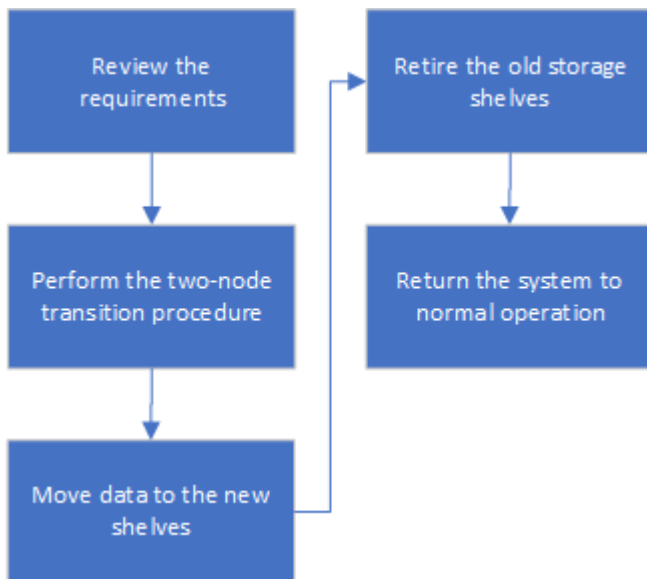
[NetApp Hardware Universe](#)

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 moving data and retiring old storage shelves

You must follow the specific workflow to ensure a successful transition.

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.



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 using the steps in [Mapping ports from the MetroCluster FC nodes to the MetroCluster IP nodes](#).
2. Prepare the MetroCluster IP controllers using the steps in [Preparing the MetroCluster IP controllers](#).
3. Verify the health of the MetroCluster FC configuration.

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

4. Gather information from the MetroCluster FC configuration.

Perform the steps in the section [Gathering information from the existing controller modules before the transition](#)

5. Remove Tiebreaker monitoring, if necessary.

Perform the steps in the section [Removing the existing configuration from the Tiebreaker or other monitoring software](#)

6. Prepare and remove the existing MetroCluster FC nodes.

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

7. Connect the new MetroCluster IP nodes.

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

8. Configure the new MetroCluster IP nodes and complete transition.

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

Migrating the root aggregates

After the transition is complete, migrate the existing root aggregates leftover from the MetroCluster FC configuration to new shelves in the MetroCluster IP configuration.

This task moves the root aggregates for node_A_1-FC and node_B_1-FC to disk shelves owned by the new MetroCluster IP controllers:

1. Assign pool 0 disks on the new local storage shelf to the controller that has the root being migrated (e.g., if the root of node_A_1-FC is being migrated, assign pool 0 disks on the new shelf to node_A_1-IP)

Note that the migration *removes and does not re-create the root mirror*, so pool 1 disks do not need to be assigned before issuing the migrate command

2. Set the privilege mode to advanced: `set priv advanced`
3. Migrate the root aggregate: `system node migrate-root -node node-name -disklist disk-id1,disk-id2,diskn -raid-type raid-type`
 - The node-name is the node to which the root aggregate is being migrated.
 - The disk-id identifies the pool 0 disks on the new shelf.
 - The raid-type is normally the same as the raid-type of the existing root aggregate.

- You can use the command `job show -idjob-id-instance` to check the migration status, where `job-id` is the value provided when the `migrate-root` command is issued. For example, if the root aggregate for `node_A_1-FC` consisted of three disks with `raid_dp`, the following command would be used to migrate root to a new shelf 11:

```
system node migrate-root -node node_A_1-IP -disklist
3.11.0,3.11.1,3.11.2 -raid-type raid_dp
```

4. Wait until the migration operation completes and the node automatically reboots.
5. Assign pool 1 disks for the root aggregate on a new shelf directly connected to the remote cluster.
6. Mirror the migrated root aggregate.
7. Wait for the root aggregate to complete resynchronising.

You can use the `storage aggregate show` command to check the sync status of the aggregates.

8. Repeat these steps for the other root aggregate.

Migrating the data aggregates

Create data aggregates on the new shelves and use `volume move` to transfer the data volumes from the old shelves to the aggregates on the new shelves.

1. Move 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](#)

Retiring shelves moved from node_A_1-FC and node_A_2-FC

You retire the old storage shelves from the original MetroCluster FC configuration. These shelves were originally owned by `node_A_1-FC` and `node_A_2-FC`.

1. Identify the aggregates on the old shelves on `cluster_B` that need to be deleted.

In this example the following data aggregates are hosted by the MetroCluster FC `cluster_B` and need to be deleted: `aggr_data_a1` and `aggr_data_a2`.



You need to perform the steps to identify, offline and delete the data aggregates on the shelves. The example is for one cluster only.

```
cluster_B::> aggr show
```

```
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
aggr0_node_A_1-FC
      349.0GB   16.83GB   95% online    1 node_A_1-IP
raid_dp,

mirrored,

normal
aggr0_node_A_2-IP
      349.0GB   16.83GB   95% online    1 node_A_2-IP
raid_dp,

mirrored,

normal
...
8 entries were displayed.

cluster_B::>
```

2. Check if the data aggregates have any MDV_aud volumes, and delete them prior to deleting the aggregates.

You must delete the MDV_aud volumes as they cannot be moved.

3. Take each of the aggregates offline, and then delete them:

- a. Take the aggregate offline: `storage aggregate offline -aggregate aggregate-name`

The following example shows the aggregate `node_B_1_aggr0` being taken offline:

```
cluster_B::> storage aggregate offline -aggregate node_B_1_aggr0

Aggregate offline successful on aggregate: node_B_1_aggr0
```

- b. Delete the aggregate: `storage aggregate delete -aggregate aggregate-name`

You can destroy the plex when prompted.

The following example shows the aggregate `node_B_1_aggr0` being deleted.

```
cluster_B::> storage aggregate delete -aggregate node_B_1_aggr0
Warning: Are you sure you want to destroy aggregate "node_B_1_aggr0"?
{y|n}: y
[Job 123] Job succeeded: DONE

cluster_B::>
```

4. After deleting all aggregates, power down, disconnect, and remove the shelves.
5. Repeat the above steps to retire the cluster_A shelves.

Completing transition

With the old controller modules removed, you can complete the transition process.

1. Complete the transition process.

Perform the steps in [Returning the system to normal operation](#)

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.