



Maintenance procedures for all MetroCluster configurations

ONTAP MetroCluster

NetApp
February 13, 2026

Table of Contents

Maintenance procedures for all MetroCluster configurations	1
Replacing a shelf nondisruptively in a stretch MetroCluster configuration	1
When to migrate root volumes to a new destination	3
Migrating root volumes within a two-node MetroCluster configuration	3
Migrating root volumes within a four-node MetroCluster configuration	3
Moving a metadata volume in MetroCluster configurations	3
Renaming a cluster in MetroCluster configurations	6
Verify the health of a MetroCluster configuration	8
Where to find additional information	10

Maintenance procedures for all MetroCluster configurations

Replacing a shelf nondisruptively in a stretch MetroCluster configuration

You can replace disk shelves without disruption in a stretch MetroCluster configuration with a fully populated disk shelf or a disk shelf chassis and transfer components from the shelf you are removing.

The disk shelf model you are installing must meet the storage system requirements specified in the [Hardware Universe](#), which includes supported shelf models, supported disk drive types, the maximum number of disk shelves in a stack, and supported ONTAP versions.

Steps

1. Properly ground yourself.
2. Identify all aggregates and volumes that have disks from the loop that contains the shelf you are replacing and make note of the affected plex name.
Either node might contain disks from the loop of the affected shelf and host aggregates or host volumes.
3. Choose one of the following two options based on the replacement scenario you are planning.
 - If you are replacing a complete disk shelf, including the shelf chassis, disks, and I/O modules (IOM), take the corresponding action as described in the table below:

Scenario	Action
The affected plex contains fewer disks from the affected shelf.	Replace the disks one-by-one on the affected shelf with spares from another shelf.  You can take the plex offline after completing the disk replacement.
The affected plex contains more disks than are in the affected shelf.	Move the plex offline and then delete the plex.
The affected plex has any disk from the affected shelf.	Move the plex offline but do not delete it.

- If you are replacing only the disk shelf chassis and no other components, perform the following steps:

- a. Offline the affected plexes from the controller where they are hosted:

```
aggregate offline
```

- b. Verify that the plexes are offline:

```
aggregate status -r
```

4. Identify the controller SAS ports to which the affected shelf loop is connected and disable the SAS ports on both site controllers:

```
storage port disable -node node_name -port SAS_port
```

The affected shelf loop is connected to both sites.

5. Wait for ONTAP to recognize that the disk is missing.

- a. Verify that the disk is missing:

```
sysconfig -a or sysconfig -r
```

6. Turn off the power switch on the disk shelf.

7. Unplug all power cords from the disk shelf.

8. Make a record of the ports from which you unplug the cables so that you can cable the new disk shelf in the same way.

9. Unplug and remove the cables connecting the disk shelf to the other disk shelves or the storage system.

10. Remove the disk shelf from the rack.

To make the disk shelf lighter and easier to maneuver, remove the power supplies and IOM. If you will be installing a disk shelf chassis, also remove the disk drives or carriers. Otherwise, avoid removing disk drives or carriers if possible because excessive handling can cause internal drive damage.

11. Install and secure the replacement disk shelf onto the support brackets and rack.

12. If you installed a disk shelf chassis, reinstall power supplies and IOM.

13. Reconfigure the stack of disk shelves by connecting all cables to the replacement disk shelf ports exactly as they were configured on the disk shelf that you removed.

14. Turn on the power to the replacement disk shelf and wait for the disk drives to spin up.

15. Change the disk shelf ID to a unique ID from 0 through 98.

16. Enable any SAS ports that you previously disabled .

- a. Wait for ONTAP to recognize that the disks are inserted.

- b. Verify that the disks are inserted:

```
sysconfig -a or sysconfig -r
```

17. If you are replacing the complete disk shelf (disk shelf chassis, disks, IOM), perform the following steps:



If you are replacing only the disk shelf chassis and no other components, go to Step 19.

- a. Determine whether disk auto assignment is enabled (on).

```
storage disk option modify -autoassign
```

Disk assignment will occur automatically.

- b. If disk auto assignment is not enabled, assign disk ownership manually.

18. Move the plexes back online:

```
aggregate online plex name
```

19. Recreate any plexes that were deleted by mirroring the aggregate.

20. Monitor the plexes as they begin resynchronizing:

```
aggregate status -r <aggregate name>
```

21. Verify that the storage system is functioning as expected:

```
system health alert show
```

When to migrate root volumes to a new destination

You might need to move root volumes to another root aggregate within a two-node or four-node MetroCluster configuration.

Migrating root volumes within a two-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a two-node MetroCluster configuration, you should refer to [How to move mroot to a new root aggregate in a 2-node Clustered MetroCluster with Switchover](#). This procedure shows you how to non-disruptively migrate the root volumes during a MetroCluster switchover operation. This procedure is slightly different than the procedure used on a four-node configuration.

Migrating root volumes within a four-node MetroCluster configuration

To migrate root volumes to a new root aggregate within a four-node MetroCluster configuration, you can use the [system node migrate-root](#) command while meeting the following requirements.

- You can use system node migrate-root to move root aggregates within a four-node MetroCluster configuration.
- All root aggregates must be mirrored.
- You can add new shelves on both sites with smaller drives to host the root aggregate.
- You must check the drive limits that the platform supports before attaching new drives.

[NetApp Hardware Universe](#)

- If you move the root aggregate to smaller drives, you need to accommodate the minimum root volume size of the platform to ensure all core files are saved.



The four-node procedure can also be applied to an eight-node configuration.

Moving a metadata volume in MetroCluster configurations

You can move a metadata volume from one aggregate to another aggregate in a MetroCluster configuration. You might want to move a metadata volume when the source aggregate is decommissioned or unmirrored, or for other reasons that make the aggregate ineligible.

- You must have cluster administrator privileges to perform this task.
- The target aggregate must be mirrored and should not be in the degraded state.
- The available space in the target aggregate must be larger than the metadata volume that you are moving.

Steps

1. Set the privilege level to advanced:

```
set -privilege advanced
```

2. Identify the metadata volume that should be moved:

```
volume show MDV_CRS*
```

```
Cluster_A:::> volume show MDV_CRS*
Vserver      Volume          Aggregate      State      Type      Size
Available    Used%
-----  -----
-----  -----
Cluster_A
      MDV_CRS_14c00d4ac9f311e7922800a0984395f1_A
          Node_A_1_aggr1
                  online      RW      10GB
  9.50GB      5%
Cluster_A
      MDV_CRS_14c00d4ac9f311e7922800a0984395f1_B
          Node_A_2_aggr1
                  online      RW      10GB
  9.50GB      5%
Cluster_A
      MDV_CRS_15035e66c9f311e7902700a098439625_A
          Node_B_1_aggr1
                  -          RW      -
  -
  -
Cluster_A
      MDV_CRS_15035e66c9f311e7902700a098439625_B
          Node_B_2_aggr1
                  -          RW      -
  -
  -
4 entries were displayed.

Cluster_A::>
```

3. Identify an eligible target aggregate:

```
metrocluster check config-replication show-aggregate-eligibility
```

The following command identifies the aggregates in cluster_A that are eligible to host metadata volumes:

```
Cluster_A::* metrocluster check config-replication show-aggregate-eligibility

Aggregate Hosted Config Replication Vols Host Addl Vols Comments
----- -----
----- 
Node_A_1_aggr0 - false Root Aggregate
Node_A_2_aggr0 - false Root Aggregate
Node_A_1_aggr1 MDV CRS 1bc7134a5ddf11e3b63f123478563412_A true -
Node_A_2_aggr1 MDV CRS 1bc7134a5ddf11e3b63f123478563412_B true -
Node_A_1_aggr2 - true
Node_A_2_aggr2 - true
Node_A_1_Aggr3 - false Unable to determine available space of aggregate
Node_A_1_aggr5 - false Unable to determine mirror configuration
Node_A_2_aggr6 - false Mirror configuration does not match requirement
Node_B_1_aggr4 - false NonLocal Aggregate
```



In the previous example, Node_A_1_aggr2 and Node_A_2_aggr2 are eligible.

4. Start the volume move operation:

```
volume move start -vserver svm_name -volume metadata_volume_name -destination-aggregate destination_aggregate_name
```

The following command moves metadata volume MDV CRS 14c00d4ac9f311e7922800a0984395f1 from aggregate Node_A_1_aggr1 to aggregate Node_A_1_aggr2:

```
Cluster_A::* volume move start -vserver svm_cluster_A -volume
MDV CRS 14c00d4ac9f311e7922800a0984395f1
-destination-aggregate aggr_cluster_A_02_01

Warning: You are about to modify the system volume
"MDV CRS 9da04864ca6011e7b82e0050568be9fe_A". This may cause
severe
    performance or stability problems. Do not proceed unless
directed to
    do so by support. Do you want to proceed? {y|n}: y
[Job 109] Job is queued: Move
"MDV CRS 9da04864ca6011e7b82e0050568be9fe_A" in Vserver
"svm_cluster_A" to aggregate "aggr_cluster_A_02_01".
Use the "volume move show -vserver svm_cluster_A -volume
MDV CRS 9da04864ca6011e7b82e0050568be9fe_A" command to view the status
of this operation.
```

5. Verify the state of the volume move operation:

```
volume move show -volume vol_constituent_name
```

6. Return to the admin privilege level:

```
set -privilege admin
```

Renaming a cluster in MetroCluster configurations

Renaming a cluster in a MetroCluster configuration involves making the changes, and then verifying on both the local and remote clusters that the change took effect correctly.

Steps

1. View the cluster names using the

```
metrocluster node show
```

command:

```
cluster_1::>* metrocluster node show
DR                         Configuration  DR
Group Cluster Node          State        Mirroring Mode
-----  -----  -----  -----
-----  -----
1      cluster_1
      node_A_1      configured  enabled   normal
      node_A_2      configured  enabled   normal
      cluster_2
      node_B_1      configured  enabled   normal
      node_B_2      configured  enabled   normal
4 entries were displayed.
```

2. Rename the cluster:

```
cluster identity modify -name new_name
```

In the following example, the cluster_1 cluster is renamed cluster_A:

```
cluster_1::>* cluster identity modify -name cluster_A
```

3. Verify on the local cluster that the renamed cluster is running normally:

```
metrocluster node show
```

In the following example, the newly renamed cluster_A is running normally:

```
cluster_A::*> metrocluster node show
DR                               Configuration  DR
Group Cluster Node              State        Mirroring Mode
-----  -----
-----  -----
1      cluster_A
      node_A_1      configured  enabled   normal
      node_A_2      configured  enabled   normal
cluster_2
      node_B_1      configured  enabled   normal
      node_B_2      configured  enabled   normal
4 entries were displayed.
```

4. Rename the remote cluster:

```
cluster peer modify-local-name -name cluster_2 -new-name cluster_B
```

In the following example, `cluster_2` is renamed `cluster_B`:

```
cluster_A::> cluster peer modify-local-name -name cluster_2 -new-name
cluster_B
```

5. Verify on the remote cluster that the local cluster was renamed and is running normally:

```
metrocluster node show
```

In the following example, the newly renamed `cluster_B` is running normally:

```
cluster_B::*> metrocluster node show
DR                               Configuration  DR
Group Cluster Node              State        Mirroring Mode
-----  -----
-----  -----
1      cluster_B
      node_B_1      configured  enabled   normal
      node_B_2      configured  enabled   normal
cluster_A
      node_A_1      configured  enabled   normal
      node_A_2      configured  enabled   normal
4 entries were displayed.
```

6. Repeat these steps for each cluster that you want to rename.

Verify the health of a MetroCluster configuration

Learn how to verify that the MetroCluster components are healthy.

About this task

- In MetroCluster IP and FC configurations, you can use the CLI to run health check commands and verify the state of the MetroCluster components.
- In MetroCluster IP configurations running ONTAP 9.8 or later, you can also use ONTAP System Manager to monitor and troubleshoot health check alerts.

Steps

Verify the health of the MetroCluster configuration depending on whether you are using the CLI or System Manager.

CLI

Use the follow steps to check the health of a MetroCluster configuration using the CLI.

Steps

1. Verify that the MetroCluster components are healthy:

```
metrocluster check run
```

```
cluster_A:::*> metrocluster check run
```

The operation runs in the background.

2. After the `metrocluster check run` operation completes, display the results:

```
metrocluster check show
```

After approximately five minutes, the following results are displayed:

```
cluster_A:::*> metrocluster check show
```

Component	Result
nodes	ok
lifs	ok
config-replication	ok
aggregates	ok
clusters	ok
connections	ok
volumes	ok
7 entries were displayed.	

3. Check the status of the running MetroCluster check operation:

```
metrocluster operation history show -job-id <id>
```

4. Verify that there are no health alerts:

```
system health alert show
```

Beginning with ONTAP 9.8, System Manager monitors the health of MetroCluster IP configurations and helps you identify and correct problems that might occur.

System Manager periodically checks the health of your MetroCluster IP configuration. When you view the MetroCluster section in the Dashboard, usually the message is "MetroCluster systems are healthy."

However, when a problem occurs, the message will show the number of events. You can click on this message and view the results of the health check for the following components:

- Node
- Network Interface
- Tier (Storage)
- Cluster
- Connection
- Volume
- Configuration Replication

The **Status** column identifies which components have problems, and the **Details** column suggests how to correct the problem.

Steps

1. In System Manager, select **Dashboard**.
2. View the message in the **MetroCluster** section:
 - a. If the message indicates that your MetroCluster configuration is healthy, and the connections between the clusters and the ONTAP Mediator are healthy (shown with check marks), then you have no problems to correct.
 - b. If the message lists the number of events, or the connections have gone down (shown with an "X"), then continue to the next step.
3. Click the message that shows the number of events.

The MetroCluster Health Report displays.

4. Troubleshoot the problems that appear in the report using the suggestions in the **Details** column.
5. When all the problems have been corrected, click **Check MetroCluster Health**.



You should perform all your troubleshooting tasks before running the check because the MetroCluster Health Check uses an intensive amount of resources.

The MetroCluster Health Check runs in the background. You can work on other tasks while you wait for it to finish.

Where to find additional information

You can learn more about configuring, operating, and monitoring a MetroCluster configuration in NetApp's extensive documentation.

Information	Subject
MetroCluster documentation	<ul style="list-style-type: none"> • All MetroCluster information
NetApp MetroCluster Solution Architecture and Design	<ul style="list-style-type: none"> • A technical overview of the MetroCluster configuration and operation. • Best practices for MetroCluster configuration.
Fabric-attached MetroCluster installation and configuration	<ul style="list-style-type: none"> • Fabric-attached MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the FC switches • Configuring the MetroCluster in ONTAP
Stretch MetroCluster installation and configuration	<ul style="list-style-type: none"> • Stretch MetroCluster architecture • Cabling the configuration • Configuring the FC-to-SAS bridges • Configuring the MetroCluster in ONTAP
MetroCluster IP installation and configuration	<ul style="list-style-type: none"> • MetroCluster IP architecture • Cabling the MetroCluster IP configuration • Configuring the MetroCluster in ONTAP
NetApp Documentation: Product Guides and Resources	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration and performance
MetroCluster Tiebreaker Software installation and configuration	<ul style="list-style-type: none"> • Monitoring the MetroCluster configuration with the MetroCluster Tiebreaker software
Copy-based transition	<ul style="list-style-type: none"> • Transitioning data from 7-Mode storage systems to clustered storage systems

Copyright information

Copyright © 2026 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—with 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.