



Configure the ONTAP Mediator service for unplanned automatic switchover

ONTAP MetroCluster

NetApp
March 29, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap-metrocluster/install-ip/concept_mediator_requirements.html on March 29, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Configure the ONTAP Mediator service for unplanned automatic switchover 1
 - Prepare to install the ONTAP Mediator service 1
 - Configure the ONTAP Mediator service from a MetroCluster IP configuration 3
 - Unconfigure the ONTAP Mediator service from the MetroCluster IP configuration 6
 - Connecting a MetroCluster configuration to a different ONTAP Mediator instance 7
 - How the ONTAP Mediator supports automatic unplanned switchover. 7

Configure the ONTAP Mediator service for unplanned automatic switchover

Prepare to install the ONTAP Mediator service

Your environment must meet certain requirements.

The following requirements apply to one disaster recovery group (DR group). Learn more about [DR groups](#).

- If you plan on updating your Linux version, do so before you install the most current ONTAP Mediator service.
- The ONTAP Mediator service and MetroCluster Tiebreaker software should not both be used with the same MetroCluster configuration.
- The ONTAP Mediator must be installed on a LINUX host at a separate location from the MetroCluster sites.

The connectivity between the ONTAP Mediator and each site must be two separate failure domains.

- The ONTAP Mediator service can support up to five MetroCluster configurations simultaneously.
- Automatic unplanned switchover is supported in ONTAP 9.7 and later.

Network requirements for using Mediator in a MetroCluster configuration

To install the ONTAP Mediator service in a MetroCluster configuration, you must make sure that the configuration meets several network requirements.

- Latency

Maximum latency of less than 75ms (RTT).

Jitter must be no more than 5ms.

- MTU

The MTU size must be at least 1400.

- Packet loss

For both Internet Control Message Protocol (ICMP) and TCP traffic, packet loss must be less than 0.01%.

- Bandwidth

The link between the Mediator service and one DR group must have at least 20Mbps of bandwidth.

- Independent connectivity

Independent connectivity between each site and the ONTAP Mediator is required. A failure in one site must not interrupt the IP connectivity between the other two unaffected sites.

Host requirements for the ONTAP Mediator in a MetroCluster configuration

You must ensure that the configuration meets several host requirements.

- ONTAP Mediator must be installed at an external site that is physically separated from the two ONTAP clusters.
- ONTAP Mediator supports a maximum number of five MetroCluster configurations.
- ONTAP Mediator does not require more than the host operating system's minimum requirements for CPU and memory (RAM).
- In addition to the host operating system's minimum requirements, at least 30GB of additional usable disk space must be available.
 - Each DR group requires up to 200MB of disk space.

Firewall requirements for ONTAP Mediator

ONTAP Mediator uses a number of ports to communicate with specific services.

If you are using a third-party firewall:

- HTTPS access must be enabled.
- It must be configured to allow access on ports 31784 and 3260.

When using the default Red Hat or CentOS firewall, the firewall is automatically configured during Mediator installation.

The following table lists the ports that you must allow in your firewall:



The iSCSI port is only required in a MetroCluster IP configuration.

Port/services	Source	Destination	Purpose
31784/tcp	ONTAP cluster management interfaces	ONTAP Mediator web server	REST API (HTTPS)
3260/tcp	ONTAP cluster (data LIF or data management LIF)	ONTAP Mediator iSCSI targets	iSCSI data connection for mailboxes

Guidelines for upgrading the ONTAP Mediator in a MetroCluster configuration

If you are upgrading the ONTAP Mediator you must meet the Linux version requirements and follow guidelines for the upgrade.

- The Mediator service can be upgraded from version from an immediately prior version to the current version.
- All Mediator versions are supported on MetroCluster IP configurations running ONTAP 9.7 or later.

[Install or upgrade the ONTAP Mediator service](#)

After the upgrade

After the Mediator and operating system upgrade is complete, you should issue the `storage iscsi-initiator show` command to confirm that the Mediator connections are up.

Configure the ONTAP Mediator service from a MetroCluster IP configuration

The ONTAP Mediator service must be configured on the ONTAP node for use in a MetroCluster IP configuration.

Before you begin

- The ONTAP Mediator must have been successfully installed on a network location that can be reached by both MetroCluster sites.

[Install or upgrade the ONTAP Mediator service](#)

- You must have the IP address of the host running the ONTAP Mediator service.
- You must have the username and password for the ONTAP Mediator service.
- All nodes of the MetroCluster IP configuration must be online.



Beginning with ONTAP 9.12.1, you can enable the MetroCluster automatic forced switchover feature in a MetroCluster IP configuration. This feature is an extension of the Mediator-assisted unplanned switchover. Before you enable this feature, review the [Risks and limitations of using MetroCluster automatic forced switchover](#).

About this task

- This task enables automatic unplanned switchover by default.
- This task can be performed on the ONTAP interface of any node in the MetroCluster IP configuration.
- A single installation of the ONTAP Mediator service can be configured with up to five MetroCluster IP configurations.

Steps

1. Add the ONTAP Mediator service to ONTAP:

```
metrocluster configuration-settings mediator add -mediator-address ip-address-of-mediator-host
```



You will be prompted for the username and password for the Mediator admin user account.

2. Verify that the automatic switchover feature is enabled:

```
metrocluster show
```

3. Verify that the Mediator is now running.

- a. Show the Mediator virtual disks:

```
storage disk show -container-type mediator
```

```
cluster_A::> storage disk show -container-type mediator
```

Container	Usable	Disk	Container
Disk	Size	Shelf Bay Type	Type
Owner			Name
NET-1.5	-	- - - VMDISK	mediator -
node_A_2			
NET-1.6	-	- - - VMDISK	mediator -
node_B_1			
NET-1.7	-	- - - VMDISK	mediator -
node_B_2			
NET-1.8	-	- - - VMDISK	mediator -
node_A_1			

b. Set the privilege mode to advanced:

```
set advanced
```

```
cluster_A::> set advanced
```

c. Display the initiators labelled as mediator:

```
storage iscsi-initiator show -label mediator
```

```

cluster_A::*> storage iscsi-initiator show -label mediator
(storage iscsi-initiator show)
+
Status
Node Type Label      Target Portal      Target Name
Admin/Op
-----
node_A_1
  mailbox
    mediator 1.1.1.1      iqn.2012-
05.local:mailbox.target.6616cd3f-9ef1-11e9-aada-
00a098ccf5d8:a05e1ffb-9ef1-11e9-8f68- 00a098cbca9e:1 up/up
node_A_2
  mailbox
    mediator 1.1.1.1      iqn.2012-
05.local:mailbox.target.6616cd3f-9ef1-11e9-aada-
00a098ccf5d8:a05e1ffb-9ef1-11e9-8f68-00a098cbca9e:1 up/up

```

d. Verify the state of the automatic unplanned switchover (AUSO) failure domain:

```
metrocluster show
```



The following example output applies to ONTAP 9.13.1 and later. For ONTAP 9.12.1 and earlier, the AUSO failure domain state should be `auso-on-cluster-disaster`.

```

cluster_A::> metrocluster show
Cluster              Entry Name              State
-----
Local: cluster_A     Configuration state configured
Mode                  normal
AUSO Failure Domain auso-on-dr-group-
disaster
Remote: cluster_B    Configuration state configured
Mode                  normal
AUSO Failure Domain auso-on-dr-group-
disaster

```

4. Optionally, configure MetroCluster automatic forced switchover.

You can only use the following command in advanced privilege level.



Before using this command, review the [Risks and limitations of using MetroCluster automatic forced switchover](#).

```
metrocluster modify -allow-auto-forced-switchover true
```

Example

```
cluster_A::*> metrocluster modify -allow-auto-forced-switchover true
```

Unconfigure the ONTAP Mediator service from the MetroCluster IP configuration

You can unconfigure the ONTAP Mediator service from the MetroCluster IP configuration.

Before you begin

You must have successfully installed and configured the ONTAP Mediator on a network location that can be reached by both MetroCluster sites.

Steps

1. Unconfigure the ONTAP Mediator service by using the following command:

```
metrocluster configuration-settings mediator remove
```

You are prompted for the user name and password for the ONTAP Mediator admin user account.



If the ONTAP Mediator is down, the `metrocluster configuration-settings mediator remove` command still prompts you to enter the user name and password for the ONTAP Mediator admin user account and removes the ONTAP Mediator service from the MetroCluster configuration.

- a. Check if there are any broken disks by using the following command:

```
disk show -broken
```

Example

```
There are no entries matching your query.
```

2. Confirm that the ONTAP Mediator service has been removed from the MetroCluster configuration by running the following commands on both clusters:

- a. `metrocluster configuration-settings mediator show`

Example

```
This table is currently empty.
```

- b. `storage iscsi-initiator show -label mediator`

Example

There are no entries matching your query.

Connecting a MetroCluster configuration to a different ONTAP Mediator instance

If you want to connect the MetroCluster nodes to a different ONTAP Mediator instance, you must unconfigure and then reconfigure the Mediator connection in the ONTAP software.

Before you begin

You need the username, password, and IP address of the new ONTAP Mediator instance.

About this task

These commands can be issued from any node in the MetroCluster configuration.

Steps

1. Remove the current ONTAP Mediator from the MetroCluster configuration:

```
metrocluster configuration-settings mediator remove
```

2. Establish the new ONTAP Mediator connection to the MetroCluster configuration:

```
metrocluster configuration-settings mediator add -mediator-address ip-address-of-mediator-host
```

How the ONTAP Mediator supports automatic unplanned switchover

The ONTAP Mediator provides mailbox LUNs to store state information about the MetroCluster IP nodes. These LUNs are co-located with the ONTAP Mediator, which runs on a Linux host physically separate from the MetroCluster sites. The MetroCluster IP nodes can use the mailbox information to monitor the state of their disaster recovery (DR) partners and implement a Mediator-assisted unplanned switchover (MAUSO) in the case of a disaster.



MAUSO is not supported in MetroCluster FC configurations.

When a node detects a site failure requiring a switchover, it takes steps to confirm that the switchover is appropriate and, if so, performs the switchover. By default, a MAUSO is initiated for the following scenarios:

- Both SyncMirror mirroring and DR mirroring of each node's nonvolatile cache is operating and the caches and mirrors are synchronized at the time of the failure.
- None of the nodes at the surviving site are in takeover state.
- If a site disaster occurs. A site disaster is a failure of *all* nodes at the same site.

A MAUSO is *not* initiated in the following shutdown scenarios:

- You initiate a shutdown. For example, when you:
 - Halt the nodes
 - Reboot the nodes

Learn about the MAUSO features available with each ONTAP 9 release.

Beginning with...	Description
ONTAP 9.13.1	<ul style="list-style-type: none">• A MAUSO is initiated if a default scenario occurs and a fan or hardware failure initiates an environmental shutdown. Examples of hardware failures include a high or low temperature, or a power supply unit, NVRAM battery, or Service Processor heartbeat failure.• The default value for the failure domain is set to "auso-on-dr-group" in a MetroCluster IP configuration. For ONTAP 9.12.1 and earlier, the default value is set to "auso-on-cluster-disaster". <p>In an eight-node MetroCluster IP configuration, "auso-on-dr-group" triggers a MAUSO either on failure of the cluster or a HA pair in one DR group. For a HA pair, both nodes must fail at the same time.</p> <p>Optionally, you can change the failure domain setting to the "auso-on-cluster-disaster" domain using the <code>metrocluster modify -auto-switchover-failure-domain auso-on-cluster-disaster</code> command to trigger a MAUSO only if there are HA node pair failures in both DR groups.</p> <ul style="list-style-type: none">• You can change the behavior to force a MAUSO even if NVRAM is not in sync at the time of the failure.
ONTAP 9.12.1	<p>You can enable the MetroCluster automatic forced switchover feature in a MetroCluster IP configuration by using the <code>metrocluster modify -allow-auto-forced -switchover true</code> command.</p> <p>Switchover upon detection of a site failure happens automatically when you enable the MetroCluster automatic forced switchover feature. You can use this feature to supplement the MetroCluster IP automatic switchover capability.</p> <p>Risks and limitations of using MetroCluster automatic forced switchover</p> <p>When you allow a MetroCluster IP configuration to operate in automatic forced switchover mode, the following known issue might lead to data loss:</p> <ul style="list-style-type: none">• The nonvolatile memory in the storage controllers is not mirrored to the remote DR partner on the partner site. <p>Caution: You might encounter scenarios that are not mentioned. NetApp is not responsible for any data corruption, data loss, or other damage that might arise from enabling the MetroCluster automatic forced switchover feature. Do not use the MetroCluster automatic forced switchover feature if the risks and limitations are not acceptable to you.</p>

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.