



What to do after an ONTAP upgrade

ONTAP 9

NetApp
March 08, 2024

This PDF was generated from https://docs.netapp.com/us-en/ontap/upgrade/task_what_to_do_after_upgrade.html on March 08, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- What to do after an ONTAP upgrade 1
 - What to do after an ONTAP upgrade 1
 - Verify your cluster after ONTAP upgrade 1
 - Verify all LIFS are on home ports after ONTAP upgrade 4
- Special configurations 5
- Update the Disk Qualification Package 16

What to do after an ONTAP upgrade

What to do after an ONTAP upgrade

After you upgrade ONTAP, there are several tasks you should perform to verify your cluster readiness.

1. [Verify your cluster.](#)

After you upgrade ONTAP, you should verify your cluster version, cluster health, and storage health. If you are using a MetroCluster FC configuration, you also need to verify that the cluster is enabled for automatic unplanned switchover.

2. [Verify that all LIFs are on home ports.](#)

During a reboot, some LIFs might have been migrated to their assigned failover ports. After you upgrade a cluster, you must enable and revert any LIFs that are not on their home ports.

3. Verify [special considerations](#) specific to your cluster.

If certain configurations exist on your cluster, you might need to perform additional steps after you upgrade.

4. [Update the Disk Qualification Package \(DQP\).](#)

The DQP is not updated as part of an ONTAP upgrade.

Verify your cluster after ONTAP upgrade

After you upgrade ONTAP, verify the cluster version, cluster health, and storage health. For MetroCluster FC configurations, also verify that the cluster is enabled for automatic unplanned switchover.

Verify cluster version

After all the HA pairs have been upgraded, you must use the version command to verify that all of the nodes are running the target release.

The cluster version is the lowest version of ONTAP running on any node in the cluster. If the cluster version is not the target ONTAP release, you can upgrade your cluster.

1. Verify that the cluster version is the target ONTAP release:

```
version
```

2. If the cluster version is not the target ONTAP release, you should verify the upgrade status of all nodes:

```
system node upgrade-revert show
```

Verify cluster health

After you upgrade a cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum.

1. Verify that the nodes in the cluster are online and are eligible to participate in the cluster:

```
cluster show
```

```
cluster1::> cluster show
Node              Health  Eligibility
-----
node0             true   true
node1             true   true
```

If any node is unhealthy or ineligible, check EMS logs for errors and take corrective action.

2. Set the privilege level to advanced:

```
set -privilege advanced
```

3. Verify the configuration details for each RDB process.
 - The relational database epoch and database epochs should match for each node.
 - The per-ring quorum master should be the same for all nodes.

Note that each ring might have a different quorum master.

To display this RDB process...	Enter this command...
Management application	cluster ring show -unitname mgmt
Volume location database	cluster ring show -unitname vldb
Virtual-Interface manager	cluster ring show -unitname vifmgr
SAN management daemon	cluster ring show -unitname bcomd

This example shows the volume location database process:

```
cluster1::*> cluster ring show -unitname vldb
```

Node	UnitName	Epoch	DB Epoch	DB Trnxs	Master	Online
node0	vldb	154	154	14847	node0	master
node1	vldb	154	154	14847	node0	secondary
node2	vldb	154	154	14847	node0	secondary
node3	vldb	154	154	14847	node0	secondary

4 entries were displayed.

- If you are operating in a SAN environment, verify that each node is in a SAN quorum:

```
cluster kernel-service show
```

```
cluster1::*> cluster kernel-service show
```

Master	Cluster	Quorum	Availability
Operational			
Node	Node	Status	Status
cluster1-01	cluster1-01	in-quorum	true
cluster1-02	cluster1-02	in-quorum	true

2 entries were displayed.

Related information

[System administration](#)

Verify automatic unplanned switchover is enabled (MetroCluster FC configurations only)

If your cluster is in a MetroCluster FC configuration, you should verify that automatic unplanned switchover is enabled after you upgrade ONTAP.

If you are using a MetroCluster IP configuration, skip this procedure.

Steps

- Check whether automatic unplanned switchover is enabled:

```
metrocluster show
```

If automatic unplanned switchover is enabled, the following statement appears in the command output:

```
AUSO Failure Domain  auso-on-cluster-disaster
```

2. If the statement does not appear, enable an automatic unplanned switchover:

```
metrocluster modify -auto-switchover-failure-domain auso-on-cluster-disaster
```

3. Verify that an automatic unplanned switchover has been enabled:

```
metrocluster show
```

Related information

[Disk and aggregate management](#)

Verify all LIFS are on home ports after ONTAP upgrade

During the reboot that occurs as part of the ONTAP upgrade process, some LIFs might be migrated from their home ports to their assigned failover ports. After an upgrade, you need to enable and revert any LIFs that are not on their home ports.

Steps

1. Display the status of all LIFs:

```
network interface show -fields home-port,curr-port
```

If **Status Admin** is "down" or **Is home** is "false" for any LIFs, continue with the next step.

2. Enable the data LIFs:

```
network interface modify {-role data} -status-admin up
```

3. Revert LIFs to their home ports:

```
network interface revert *
```

4. Verify that all LIFs are in their home ports:

```
network interface show
```

This example shows that all LIFs for SVM vs0 are on their home ports.

```
cluster1::> network interface show -vserver vs0
```

Vserver	Logical Interface	Status Admin/Oper	Network Address/Mask	Current Node	Current Port	Is Home
vs0						
	data001	up/up	192.0.2.120/24	node0	e0e	true
	data002	up/up	192.0.2.121/24	node0	e0f	true
	data003	up/up	192.0.2.122/24	node0	e2a	true
	data004	up/up	192.0.2.123/24	node0	e2b	true
	data005	up/up	192.0.2.124/24	node1	e0e	true
	data006	up/up	192.0.2.125/24	node1	e0f	true
	data007	up/up	192.0.2.126/24	node1	e2a	true
	data008	up/up	192.0.2.127/24	node1	e2b	true

8 entries were displayed.

Special configurations

Special considerations after an ONTAP upgrade

If your cluster is configured with any of the following features you might need to perform additional steps after you upgrade your ONTAP software.

Ask yourself...	If your answer is yes, then do this...
Did I upgrade from ONTAP 9.7 or earlier to ONTAP 9.8 or later?	Verify your network configuration Remove the EMS LIF service from network service policies that do not provide reachability to the EMS destination
Is my cluster in a MetroCluster configuration?	Verify your networking and storage status
Do I have a SAN configuration?	Verify your SAN configuration
Did I upgrade from ONTAP 9.3 or earlier, and am using NetApp Storage Encryption?	Reconfigure KMIP server connections
Do I have load-sharing mirrors?	Relocate moved load-sharing mirror source volumes
Does my cluster have defined SnapMirror relationships?	Resume SnapMirror operations
Do I have user accounts for Service Processor (SP) access that were created prior to ONTAP 9.9.1?	Verify the change in accounts that can access the Service Processor

Verify your networking configuration after an ONTAP upgrade from ONTAP 9.7x or earlier

After you upgrade from ONTAP 9.7x or earlier to ONTAP 9.8 or later, you should verify your network configuration. After the upgrade, ONTAP automatically monitors layer 2


reachability.

Step

1. Verify each port has reachability to its expected broadcast domain:

```
network port reachability show -detail
```

The command output contains reachability results. Use the following decision tree and table to understand the reachability results (reachability-status) and determine what, if anything, to do next.



reachability-status	Description
---------------------	-------------

ok	<p>The port has layer 2 reachability to its assigned broadcast domain.</p> <p>If the reachability-status is "ok", but there are "unexpected ports", consider merging one or more broadcast domains. For more information, see Merge broadcast domains.</p> <p>If the reachability-status is "ok", but there are "unreachable ports", consider splitting one or more broadcast domains. For more information, see Split broadcast domains.</p> <p>If the reachability-status is "ok", and there are no unexpected or unreachable ports, your configuration is correct.</p>
misconfigured-reachability	<p>The port does not have layer 2 reachability to its assigned broadcast domain; however, the port does have layer 2 reachability to a different broadcast domain.</p> <p>You can repair the port reachability. When you run the following command, the system will assign the port to the broadcast domain to which it has reachability:</p> <pre>network port reachability repair -node -port</pre> <p>For more information, see Repair port reachability.</p>
no-reachability	<p>The port does not have layer 2 reachability to any existing broadcast domain.</p> <p>You can repair the port reachability. When you run the following command, the system will assign the port to a new automatically created broadcast domain in the Default IPspace:</p> <pre>network port reachability repair -node -port</pre> <p>For more information, see Repair port reachability.</p>
multi-domain-reachability	<p>The port has layer 2 reachability to its assigned broadcast domain; however, it also has layer 2 reachability to at least one other broadcast domain.</p> <p>Examine the physical connectivity and switch configuration to determine if it is incorrect or if the port's assigned broadcast domain needs to be merged with one or more broadcast domains.</p> <p>For more information, see Merge broadcast domains or Repair port reachability.</p>
unknown	<p>If the reachability-status is "unknown", then wait a few minutes and try the command again.</p>

After you repair a port, you need to check for and resolve displaced LIFs and VLANs. If the port was part of an interface group, you also need to understand what happened to that interface group. For more information, see [Repair port reachability](#).

Remove EMS LIF service from network service policies

If you have Event Management System (EMS) messages set up before you upgrade from ONTAP 9.7 or earlier to ONTAP 9.8 or later , after the upgrade, your EMS messages

might not be delivered.

During the upgrade, management-ems, which is the EMS LIF service, is added to all existing service policies. This allows EMS messages to be sent from any of the LIFs associated with any of the service policies. If the selected LIF does not have reachability to the event notification destination, the message is not delivered.

To prevent this, after the upgrade, you should remove the EMS LIF service from the network service policies that do not provide reachability to the destination.

Steps

1. Identify the LIFs and associated network service policies through which EMS messages can be sent:

```
network interface show -fields service-policy -services management-ems
```

vserver	lif	service-policy
cluster-1	cluster_mgmt	
		default-management
cluster-1	node1-mgmt	
		default-management
cluster-1	node2-mgmt	
		default-management
cluster-1	inter_cluster	
		default-intercluster

4 entries were displayed.

2. Check each LIF for connectivity to the EMS destination:

```
network ping -lif lif_name -vserver svm_name -destination  
destination_address
```

Perform this on each node.

Examples

```
cluster-1::> network ping -lif node1-mgmt -vserver cluster-1  
-destination 10.10.10.10  
10.10.10.10 is alive  
  
cluster-1::> network ping -lif inter_cluster -vserver cluster-1  
-destination 10.10.10.10  
no answer from 10.10.10.10
```

3. Enter advanced privilege level:

```
set advanced
```

4. For the LIFs that do not have reachability, remove the management-ems LIF service from the corresponding service policies:

```
network interface service-policy remove-service -vserver svm_name  
-policy service_policy_name -service management-ems
```

5. Verify that the management-ems LIF is now only associated with the LIFs that provide reachability to the EMS destination:

```
network interface show -fields service-policy -services management-ems
```

Related Links

[LIFs and service policies in ONTAP 9.6 and later](#)

Verify networking and storage status for MetroCluster configurations after an ONTAP upgrade

After you upgrade an ONTAP cluster in a MetroCluster configuration, you should verify the status of the LIFs, aggregates, and volumes for each cluster.

1. Verify the LIF status:

```
network interface show
```

In normal operation, LIFs for source SVMs must have an admin status of up and be located on their home nodes. LIFs for destination SVMs are not required to be up or located on their home nodes. In switchover, all LIFs have an admin status of up, but they do not need to be located on their home nodes.

```

cluster1::> network interface show

```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	
Cluster					
	cluster1-a1_clus1	up/up	192.0.2.1/24	cluster1-01	e2a
true					
	cluster1-a1_clus2	up/up	192.0.2.2/24	cluster1-01	e2b
true					
cluster1-01					
	clus_mgmt	up/up	198.51.100.1/24	cluster1-01	e3a
true					
	cluster1-a1_inet4_intercluster1	up/up	198.51.100.2/24	cluster1-01	e3c
true					
	...				

```

27 entries were displayed.

```

2. Verify the state of the aggregates:

```
storage aggregate show -state !online
```

This command displays any aggregates that are *not* online. In normal operation, all aggregates located at the local site must be online. However, if the MetroCluster configuration is in switchover, root aggregates at the disaster recovery site are permitted to be offline.

This example shows a cluster in normal operation:

```

cluster1::> storage aggregate show -state !online
There are no entries matching your query.

```

This example shows a cluster in switchover, in which the root aggregates at the disaster recovery site are

offline:

```
cluster1::> storage aggregate show -state !online
Aggregate      Size Available Used% State  #Vols  Nodes      RAID
Status
-----
-----
aggr0_b1
      0B      0B    0% offline    0 cluster2-01
raid_dp,
mirror
degraded
aggr0_b2
      0B      0B    0% offline    0 cluster2-02
raid_dp,
mirror
degraded
2 entries were displayed.
```

3. Verify the state of the volumes:

```
volume show -state !online
```

This command displays any volumes that are *not* online.

If the MetroCluster configuration is in normal operation (it is not in switchover state), the output should show all volumes owned by the cluster's secondary SVMs (those with the SVM name appended with "-mc").

Those volumes come online only in the event of a switchover.

This example shows a cluster in normal operation, in which the volumes at the disaster recovery site are not online.

```
cluster1::> volume show -state !online
(volume show)
Vserver   Volume      Aggregate    State    Type    Size
Available Used%
-----
vs2-mc    vol1         aggr1_b1     -        RW      -
-         -
vs2-mc    root_vs2     aggr0_b1     -        RW      -
-         -
vs2-mc    vol2         aggr1_b1     -        RW      -
-         -
vs2-mc    vol3         aggr1_b1     -        RW      -
-         -
vs2-mc    vol4         aggr1_b1     -        RW      -
-         -
5 entries were displayed.
```

4. Verify that there are no inconsistent volumes:

```
volume show -is-inconsistent true
```

See the Knowledge Base article [Volume Showing WAFL Inconsistent](#) on how to address the inconsistent volumes.

Verify the SAN configuration after an upgrade

After an ONTAP upgrade, in a SAN environment, you should verify that each initiator that was connected to a LIF before the upgrade has successfully reconnected to the LIF.

1. Verify that each initiator is connected to the correct LIF.

You should compare the list of initiators to the list you made during the upgrade preparation.

For...	Enter...
iSCSI	<pre>iscsi initiator show -fields igroup,initiator-name,tpgroup</pre>
FC	<pre>fcp initiator show -fields igroup,wwpn,lif</pre>

Reconfigure KMIP server connections after an upgrade from ONTAP 9.2 or earlier

After you upgrade from ONTAP 9.2 or earlier to ONTAP 9.3 or later, you need to reconfigure any external key management (KMIP) server connections.

Steps

1. Configure the key manager connectivity:

```
security key-manager setup
```

2. Add your KMIP servers:

```
security key-manager add -address key_management_server_ip_address
```

3. Verify that KMIP servers are connected:

```
security key-manager show -status
```

4. Query the key servers:

```
security key-manager query
```

5. Create a new authentication key and passphrase:

```
security key-manager create-key -prompt-for-key true
```

The passphrase must have a minimum of 32 characters.

6. Query the new authentication key:

```
security key-manager query
```

7. Assign the new authentication key to your self-encrypting disks (SEDs):

```
storage encryption disk modify -disk disk_ID -data-key-id key_ID
```



Make sure you are using the new authentication key from your query.

8. If needed, assign a FIPS key to the SEDs:


```
storage encryption disk modify -disk disk_id -fips-key-id  
fips_authentication_key_id
```

If your security setup requires you to use different keys for data authentication and FIPS 140-2 authentication, you should create a separate key for each. If that is not the case, you can use the same authentication key for FIPS compliance that you use for data access.

Relocate moved load-sharing mirror source volumes after an ONTAP upgrade

After you upgrade ONTAP, you need to move load-sharing mirror source volumes back to their pre-upgrade locations.

Steps

1. Identify the location to which you are moving the load-sharing mirror source volume by using the record you created before moving the load-sharing mirror source volume.
2. Move the load-sharing mirror source volume back to its original location:

```
volume move start
```

Resume SnapMirror operations after an ONTAP upgrade

After an ONTAP upgrade, you need to resume any SnapMirror relationships that were suspended.

Steps

1. Resume transfers for each SnapMirror relationship that was previously quiesced:

```
snapmirror resume *
```

This command resumes the transfers for all quiesced SnapMirror relationships.

2. Verify that the SnapMirror operations have resumed:

```
snapmirror show
```

```
cluster1::> snapmirror show
```

Source	Destination	Mirror	Relationship	Total		
Last						
Path	Type	Path	State	Status	Progress	Healthy
Updated						
-----	----	-----	-----	-----	-----	-----

cluster1-vs1:dp_src1						
	DP	cluster1-vs2:dp_dst1				
			Snapmirrored			
			Idle		-	true -
cluster1-vs1:xdp_src1						
	XDP	cluster1-vs2:xdp_dst1				
			Snapmirrored			
			Idle		-	true -
cluster1://cluster1-vs1/ls_src1						
	LS	cluster1://cluster1-vs1/ls_mr1				
			Snapmirrored			
			Idle		-	true -
		cluster1://cluster1-vs1/ls_mr2				
			Snapmirrored			
			Idle		-	true -

4 entries were displayed.

For each SnapMirror relationship, verify that the Relationship Status is **Idle**. If the status is **Transferring**, wait for the SnapMirror transfer to complete, and then reenter the command to verify that the status has changed to **Idle**.

For each SnapMirror relationship that is configured to run on a schedule, you should verify that the first scheduled SnapMirror transfer completes successfully.

Change in user accounts that can access the Service Processor

If you created user accounts in ONTAP 9.8 or earlier that can access the Service Processor (SP) with a non-admin role and you upgrade to ONTAP 9.9.1 or later, any non-admin value in the `-role` parameter is modified to `admin`.

For more information, see [Accounts that can access the SP](#).

Update the Disk Qualification Package

After you upgrade your ONTAP software, you should download and install the ONTAP Disk Qualification Package (DQP). The DQP is not updated as part of an ONTAP upgrade.

The DQP contains the proper parameters for ONTAP interaction with all newly qualified drives. If your version of the DQP does not contain information for a newly qualified drive, ONTAP will not have the information to properly configure the drive.

It is best practice to update the DQP every quarter. You should also update the DQP for the following reasons:

- Whenever you add a new drive type or size to a node in your cluster

For example, if you already have 1-TB drives and add 2-TB drives, you need to check for the latest DQP update.

- Whenever you update the disk firmware
- Whenever newer disk firmware or DQP files are available

Related information

- [NetApp Downloads: Disk Qualification Package](#)
- [NetApp Downloads: Disk Drive Firmware](#)

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.