



What to do after an ONTAP revert

ONTAP 9

NetApp
December 21, 2024

Table of Contents

- What to do after an ONTAP revert 1
 - Verify cluster and storage health after an ONTAP revert 1
 - Enable automatic switchover for MetroCluster configurations after an ONTAP revert 4
 - Enable and revert LIFs to home ports after an ONTAP revert 5
 - Enable Snapshot copy policies after an ONTAP revert 7
 - Verify IPv6 firewall entries after an ONTAP revert 8
 - Verify user accounts that can access the Service Processor after reverting to ONTAP 9.8 9

What to do after an ONTAP revert

Verify cluster and storage health after an ONTAP revert

After you revert an ONTAP cluster, you should verify that the nodes are healthy and eligible to participate in the cluster, and that the cluster is in quorum. You should also verify the status of your disks, aggregates, and volumes.

Verify cluster health

Steps

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

```
cluster show
```

In this example, the cluster is healthy and all nodes are eligible to participate in the cluster.

```
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
```

Enter *y* to continue.

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	<pre>cluster ring show -unitname mgmt</pre>

To display this RDB process...	Enter this command...
Volume location database	<code>cluster ring show -unitname vldb</code>
Virtual-Interface manager	<code>cluster ring show -unitname vifmgr</code>
SAN management daemon	<code>cluster ring show -unitname bcomd</code>

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.
```

4. Return to the admin privilege level:

```
set -privilege admin
```

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

```
event log show -severity informational -message-name scsiblade.*
```

The most recent scsiblade event message for each node should indicate that the scsi-blade is in quorum.

```
cluster1::*> event log show -severity informational -message-name
scsiblade.*
```

Time	Node	Severity	Event
MM/DD/YYYY TIME	node0	INFORMATIONAL	scsiblade.in.quorum: The scsi-blade ...
MM/DD/YYYY TIME	node1	INFORMATIONAL	scsiblade.in.quorum: The scsi-blade ...

Related information

[System administration](#)

Verify storage health

After you revert or downgrade a cluster, you should verify the status of your disks, aggregates, and volumes.

Steps

1. Verify disk status:

To check for...	Do this...
Broken disks	<ol style="list-style-type: none"> a. Display any broken disks: <div data-bbox="889 1014 1487 1152" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <pre>storage disk show -state broken</pre> </div> b. Remove or replace any broken disks.
Disks undergoing maintenance or reconstruction	<ol style="list-style-type: none"> a. Display any disks in maintenance, pending, or reconstructing states: <div data-bbox="889 1367 1487 1545" style="border: 1px solid #ccc; padding: 10px; margin: 10px 0;"> <pre>storage disk show -state maintenance pending reconstruc ting</pre> </div> b. Wait for the maintenance or reconstruction operation to finish before proceeding.

2. Verify that all aggregates are online by displaying the state of physical and logical storage, including storage aggregates:

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

This command displays the aggregates that are *not* online. All aggregates must be online before and after performing a major upgrade or reversion.

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

3. Verify that all volumes are online by displaying any volumes that are *not* online:

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

All volumes must be online before and after performing a major upgrade or reversion.

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

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.

Related information

[Disk and aggregate management](#)

Verify client access (SMB and NFS)

For the configured protocols, test access from SMB and NFS clients to verify that the cluster is accessible.

Enable automatic switchover for MetroCluster configurations after an ONTAP revert

After reverting an ONTAP MetroCluster configuration, you must enable automatic unplanned switchover to ensure that the MetroCluster configuration is fully operational.

Steps

1. Enable automatic unplanned switchover:

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

2. Validate the MetroCluster configuration:

```
metrocluster check run
```

Enable and revert LIFs to home ports after an ONTAP revert

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

The network interface revert command reverts a LIF that is not currently on its home port back to its home port, provided that the home port is operational. A LIF's home port is specified when the LIF is created; you can determine the home port for a LIF by using the network interface show command.

Steps

1. Display the status of all LIFs:

```
network interface show
```

This example displays the status of all LIFs for a storage virtual machine (SVM).

```

cluster1::> network interface show -vserver vs0
      Logical      Status      Network      Current
Current Is
Vserver  Interface  Admin/Oper  Address/Mask  Node      Port
Home
-----
vs0
      data001    down/down  192.0.2.120/24  node0    e0e
true
      data002    down/down  192.0.2.121/24  node0    e0f
true
      data003    down/down  192.0.2.122/24  node0    e2a
true
      data004    down/down  192.0.2.123/24  node0    e2b
true
      data005    down/down  192.0.2.124/24  node0    e0e
false
      data006    down/down  192.0.2.125/24  node0    e0f
false
      data007    down/down  192.0.2.126/24  node0    e2a
false
      data008    down/down  192.0.2.127/24  node0    e2b
false
8 entries were displayed.

```

If any LIFs appear with a Status Admin status of down or with an Is home status of false, 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
      Logical      Status      Network      Current
Current Is
Vserver      Interface  Admin/Oper  Address/Mask  Node      Port
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.

```

Enable Snapshot copy policies after an ONTAP revert

After reverting to an earlier version of ONTAP, you must enable Snapshot copy policies to start creating Snapshot copies again.

You are reenabling the Snapshot schedules that you disabled before you reverted to an earlier version of ONTAP.

Steps

1. Enable Snapshot copy policies for all data SVMs:

```
volume snapshot policy modify -vserver * -enabled true
```

```
snapshot policy modify pg-rpo-hourly -enable true
```

2. For each node, enable the Snapshot copy policy of the root volume:

```
run -node <node_name> vol options <volume_name> nosnap off
```

Verify IPv6 firewall entries after an ONTAP revert

A reversion from any version of ONTAP 9 might result in missing default IPv6 firewall entries for some services in firewall policies. You need to verify that the required firewall entries have been restored to your system.

Steps

1. Verify that all firewall policies are correct by comparing them to the default policies:

```
system services firewall policy show
```

The following example shows the default policies:

```
cluster1::*> system services firewall policy show
Policy          Service      Action IP-List
-----
cluster
                dns         allow  0.0.0.0/0
                http        allow  0.0.0.0/0
                https       allow  0.0.0.0/0
                ndmp        allow  0.0.0.0/0
                ntp         allow  0.0.0.0/0
                rsh         allow  0.0.0.0/0
                snmp        allow  0.0.0.0/0
                ssh         allow  0.0.0.0/0
                telnet      allow  0.0.0.0/0

data
                dns         allow  0.0.0.0/0, ::/0
                http        deny   0.0.0.0/0, ::/0
                https       deny   0.0.0.0/0, ::/0
                ndmp        allow  0.0.0.0/0, ::/0
                ntp         deny   0.0.0.0/0, ::/0
                rsh         deny   0.0.0.0/0, ::/0

.
.
.
```

2. Manually add any missing default IPv6 firewall entries by creating a new firewall policy:

```
system services firewall policy create -policy <policy_name> -service  
ssh -action allow -ip-list <ip_list>
```

3. Apply the new policy to the LIF to allow access to a network service:

```
network interface modify -vserve <svm_name> -lif <lif_name> -firewall  
-policy <policy_name>
```

Verify user accounts that can access the Service Processor after reverting to ONTAP 9.8

In ONTAP 9.9.1 and later the the `-role` parameter for user accounts is changed to `admin`. If you created user accounts on ONTAP 9.8 or earlier, upgraded to ONTAP 9.9.1 or later and then reverted back to ONTAP 9.8, the `-role` parameter is restored to its original value. You should verify that the modified values are acceptable.

During revert, if the role for an SP user has been deleted, the "rbac.spuser.role.notfound" EMS message will be logged.

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

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.