NetApp

Special configurations

ONTAP 9

NetApp
March 08, 2024

Table of Contents

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 polices that do not provide reachabiity to the EMS destination
Is my cluster in a 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
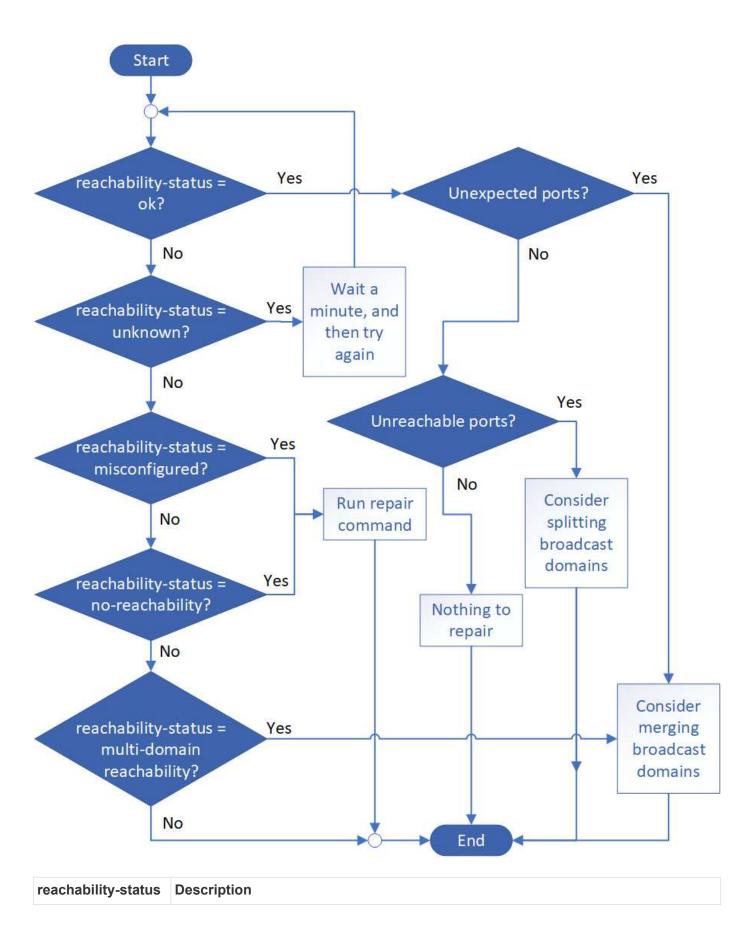
After you upgrade from ONTAP 9.7x or earlier to ONTAP 9.8 or later, you should verify your network configuration. After the upgrade, ONTAP automatically monitors layer 2 reachability.

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.

```mermaid
flowchart
    Start([Start])
```

Start

reachability-status = ok?
— Yes → Unexpected ports?
— No → reachability-status = unknown?

Unexpected ports?
— Yes → Consider merging broadcast domains
— No → Unreachable ports?

reachability-status = unknown?
— Yes → Wait a minute, and then try again
— No → reachability-status = misconfigured?

Unreachable ports?
— Yes → Consider splitting broadcast domains
— No → Nothing to repair

reachability-status = misconfigured?
— Yes → Run repair command
— No → reachability-status = no-reachability?

reachability-status = no-reachability?
— Yes → Run repair command
— No → reachability-status = multi-domain reachability?

reachability-status = multi-domain reachability?
— Yes → End
— No → End

reachability-status	Description

| ok | The port has layer 2 reachability to its assigned broadcast domain.

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.

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.

If the reachability-status is "ok", and there are no unexpected or unreachable ports, your configuration is correct. |
|---|---|
| misconfigured-reachability | 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.

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:

`network port reachability repair -node -port`

For more information, see Repair port reachability. |
| no-reachability | The port does not have layer 2 reachability to any existing broadcast domain.

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:

`network port reachability repair -node -port`

For more information, see Repair port reachability. |
| multi-domain-reachability | 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.

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.

For more information, see Merge broadcast domains or Repair port reachability. |
| unknown | If the reachability-status is "unknown", then wait a few minutes and try the command again. |

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 polices. This allows EMS messages to be sent from any of the LIFs associated with any of the service polices. 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 polices that do not provide reachability to the destination.

Steps

1. Identify the LIFs and associated network service polices 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 polices:

```
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 polices 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
            Logical     Status     Network               Current
Current Is
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.