Configure hosts with FCP and iSCSI

ONTAP SAN Host Utilities

NetApp
January 30, 2026

Table of Contents

Configure hosts with FCP and iSCSI

Overview

You can configure certain SAN hosts for FCP or iSCSI with ONTAP as the target. First you install the relevant operating system host utility package, which includes the SAN tool kit. You then verify the multipath settings for the NetApp ONTAP LUNs.

AIX and PowerVM/VIOS

Configure AIX 7.3/VIOS 4.x for FCP and iSCSI with ONTAP storage

The AIX Host Utilities software provides management and diagnostic tools for AIX hosts that are connected to ONTAP storage. When you install the AIX Host Utilities on an AIX 7.3/VIOS 4.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability. If your configuration doesn't support SAN booting, you can use a local boot.

SAN boot

SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for an AIX/PowerVM host host. You can set up a SAN boot LUN to work in a AIX Multipath I/O (MPIO) environment that is using the FC protocol and running AIX Host Utilities with either the FC or FCoE protocol. The method you use to create a SAN boot LUN and install a new OS image in an AIX MPIO environment depends on the protocol that you are using.

Steps

1. Use the Interoperability Matrix Tool to verify that your AIX OS, protocol, and ONTAP version support SAN booting.
2. Follow the best practices for setting up a SAN boot in the vendor documentation.

Local boot

Perform a local boot by installing the AIX OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install the AIX Host Utilities

NetApp strongly recommends installing the AIX Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data. The MPIO package from the Host Utilities provides MPIO support for AIX and VIOS.

Installing the AIX Host Utilities provides additional timeout settings on your AIX host.

Install AIX Host Utilities 8.0.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with AIX and PowerVM to manage ONTAP LUNs.

Multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues with the remaining paths. The AIX and PowerVM environments of the Host Utilities use the AIXs native multipathing solution (MPIO).

The Path Control Module (PCM) is responsible for controlling multiple paths for an AIX host. The PCM is a storage vendor supplied code that handles path management and is installed and enabled during the Host Utilities installation.

To ensure that multipathing is configured correctly for your host, verify that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that "MPIO NetApp" is available. "MPIO NetApp" is loaded during the AIX Host Utilities installation and becomes available after you reboot the host.

   ```
   lsdev -Cc disk
   ```

   Example output

   ```
   hdisk1 Available 00-00-02 MPIO NetApp FCP Default PCM Disk
   ```

2. The AIX Host Utilities load the following parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Environment	Value for AIX	Note
algorithm	MPIO	round_robin	Set by Host Utilities
hcheck_cmd	MPIO	inquiry	Set by Host Utilities
hcheck_interval	MPIO	30	Set by Host Utilities
hcheck_mode	MPIO	nonactive	Set by Host Utilities
lun_reset_spt	MPIO / non-MPIO	yes	Set by Host Utilities
max_transfer	MPIO / non-MPIO	FC LUNs: 0x100000 bytes	Set by Host Utilities
qfull_dly	MPIO / non-MPIO	2-second delay	Set by Host Utilities
queue_depth	MPIO / non-MPIO	64	Set by Host Utilities
reserve_policy	MPIO / non-MPIO	no_reserve	Set by Host Utilities
re_timeout (disk)	MPIO / non-MPIO	30 seconds	Uses OS Default values
dyntrk	MPIO / non-MPIO	Yes	Uses OS Default values
fc_err_recov	MPIO / non-MPIO	Fast_fail	Uses OS Default values
q_type	MPIO / non-MPIO	simple	Uses OS Default values
num_cmd_elems	MPIO / non-MPIO	1024 for AIX 3072 for VIOS	FC EN1B, FC EN1C
num_cmd_elems	MPIO / non-MPIO	1024 for AIX	FC EN0G

3. Configure the following settings to optimize I/O operations for FC.

Parameter	AIX versions	AIX OS default value	NetApp recommended value
rw_timeout (disk)	AIX 7.3TL3	NPIV:30 seconds, vSCSI:45 seconds	NPIV:30 seconds, vSCSI:120 seconds
	AIX 7.2TL5	NPIV:30 seconds, vSCSI:45 seconds	NPIV:30 seconds, vSCSI:120 seconds
	VIOS 3.1	30 seconds	30 seconds
	VIOS 4.1	30 seconds	30 seconds

4. Configure the following settings to optimize I/O operations for iSCSI.

Parameter	AIX versions	AIX OS default value	NetApp recommended value
rw_timeout (disk)	AIX 7.3TL3	vSCSI:45 seconds	vSCSI:120 seconds
	AIX 7.2TL5	vSCSI:45 seconds	vSCSI:120 seconds
	VIOS 3.1	120 seconds	30 seconds
	VIOS 4.1	120 seconds	30 seconds
	All AIX7.2 and AIX 7.3 standalone	120 seconds	30 seconds
isw_err_recov (iscsi0)	All AIX7.2 and AIX 7.3 standalone	delayed_fail	fast_fail

5. If your storage configuration includes MetroCluster or SnapMirror active sync, change the default settings:

MetroCluster

By default, the AIX OS enforces a shorter I/O timeout when there are no available paths to a LUN. This might occur in configurations that include single-switch SAN fabric and in MetroCluster configurations that experience unplanned failovers. For additional information and recommended changes to default settings, see the Knowledge Base article What are AIX Host support considerations in a MetroCluster configuration?.

SnapMirror active sync

Beginning with ONTAP 9.11.1, SnapMirror active sync is supported for an AIX host. The primary cluster in an AIX configuration is the "active" cluster.

In an AIX configuration, failovers are disruptive. With each failover, you need to perform a re-scan on the host for I/O operations to resume.

Refer to the Knowledge Base article How to configure an AIX host for SnapMirror active sync.

6. Verify the parameter settings and that multiple paths are listed for an ONTAP LUN:

```
lsmpio
```

In the following example for an AFF or FAS system, the PCM is listed for NetApp.

Show example

```
# lsmpio -l hdisk1
name    path_id  status   path_status  parent  connection

hdisk1  0        Enabled  Non          fscsi6
203200a098ba7afe,5b000000000000
hdisk1  1        Enabled  Non          fscsi8
203100a098ba7afe,5b000000000000
hdisk1  2        Enabled  Sel,Opt      fscsi6
203000a098ba7afe,5b000000000000
hdisk1  3        Enabled  Sel,Opt      fscsi8
203800a098ba7afe,5b000000000000
#
lsattr -El hdisk1
PCM             PCM/friend/NetAppDefaultPCM Path Control Module
False
PR_key_value    0x6d0000000002              Persistant Reserve Key
Value            True
algorithm       round_robin                 Algorithm
True
clr_q           no                          Device CLEARS its Queue
on error        True
dist_err_pcnt   0                           Distributed Error Sample
Time           True
dist_tw_width   50                          Distributed Error Sample
Time           True
hcheck_cmd      inquiry                     Health Check Command
True
hcheck_interval 30                          Health Check Interval
True
hcheck_mode     nonactive                   Health Check Mode
True
location                                    Location Label
True
lun_id          0x5b000000000000            Logical Unit Number ID
False
lun_reset_spt   yes                         LUN Level Reset
True
max_transfer    0x100000                    Maximum TRANSFER Size
True
node_name       0x204800a098ba7afe          FC Node Name
False
pvid            none                        Physical volume
identifier              False
q_err           yes                         Use QERR bit
```

```
True
q_type          simple                          Queuing TYPE
True
qfull_dly       2                               Delay in seconds for
SCSI TASK SET FULL True
queue_depth     64                              Queue DEPTH
True
reassign_to     120                             REASSIGN time out value
True
reserve_policy  PR_shared                       Reserve Policy
True
rw_timeout      30                              READ/WRITE time out
value               True
scsi_id         0xec409                         SCSI ID
False
start_timeout   60                              START unit time out
value               True
timeout_policy  fail_path                       Active/Passive Disk Path
Control Module True
ww_name         0x203200a098ba7afe              FC World Wide Name
False
```

7. Verify the path status for ONTAP LUNs:

```
sanlun lun show
```

The following example outputs show the correct path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configurations

An ASA configuration optimizes all paths to a given LUN, keeping them active ("primary"). This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# sanlun lun show -p |grep -p hdisk78
                    ONTAP Path:
vs_aix_clus:/vol/chataix_205p2_vol_en_1_7/jfs_205p2_lun_en
                           LUN: 37
                      LUN Size: 15g
                   Host Device: hdisk78
                          Mode: C
            Multipath Provider: AIX Native
        Multipathing Algorithm: round_robin
------ ------- ------ ------- --------- ----------
host   vserver  AIX                      AIX MPIO
path   path     MPIO   host    vserver     path
state  type     path   adapter LIF       priority
------ ------- ------ ------- --------- ----------
up     primary  path0  fcs0    fc_aix_1     1
up     primary  path1  fcs0    fc_aix_2     1
up     primary  path2  fcs1    fc_aix_3     1
up     primary  path3  fcs1    fc_aix_4     1
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the correct output for an ONTAP LUN with two Active/Optimized ("primary") paths and two Active/Non-Optimized ("secondary") paths:

Show example

```
# sanlun lun show -p |grep -p hdisk78
                    ONTAP Path:
vs_aix_clus:/vol/chataix_205p2_vol_en_1_7/jfs_205p2_lun_en
                           LUN: 37
                      LUN Size: 15g
                   Host Device: hdisk78
                          Mode: C
              Multipath Provider: AIX Native
        Multipathing Algorithm: round_robin
------- ---------- ------ ------- ---------- ----------
host    vserver    AIX                        AIX MPIO
path    path       MPIO   host    vserver         path
state   type       path   adapter LIF         priority
------- ---------- ------ ------- ---------- ----------
up      secondary  path0  fcs0    fc_aix_1         1
up      primary    path1  fcs0    fc_aix_2         1
up      primary    path2  fcs1    fc_aix_3         1
up      secondary  path3  fcs1    fc_aix_4         1
```

Step 4: Review the known issues

There are no known issues.

What's next?

Learn about using the AIX Host Utilities tool.

Configure AIX 7.2/VIOS 3.1 for FCP and iSCSI with ONTAP storage

The AIX Host Utilities software provides management and diagnostic tools for AIX hosts that are connected to ONTAP storage. When you install the AIX Host Utilities on an AIX 7.2 and/or PowerVM (VIOS 3.1) host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability. If your configuration doesn't support SAN booting, you can use a local boot.

SAN boot

SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for an AIX/PowerVM host host. You can set up a SAN boot LUN to work in a AIX Multipath I/O (MPIO) environment that is using the FC protocol and running AIX Host Utilities with either the FC or FCoE protocol. The method you use to create a SAN boot LUN and install a new OS image in an AIX MPIO environment depends on the protocol that you are using.

Steps

1. Use the Interoperability Matrix Tool to verify that your AIX OS, protocol, and ONTAP version support SAN booting.
2. Follow the best practices for setting up a SAN boot in the vendor documentation.

Local boot

Perform a local boot by installing the AIX OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install the AIX Host Utilities

NetApp strongly recommends installing the AIX Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data. The MPIO package from the Host Utilities provides MPIO support for AIX and VIOS.

Installing the AIX Host Utilities provides additional timeout settings on your AIX host.

Install AIX Host Utilities 6.1.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with AIX 7.2 and PowerVM to manage ONTAP LUNs.

Multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues with the remaining paths. The AIX and PowerVM environments of the Host Utilities use the AIXs native multipathing solution (MPIO).

The Path Control Module (PCM) is responsible for controlling multiple paths for an AIX host. The PCM is a storage vendor supplied code that handles path management and is installed and enabled during the Host Utilities installation.

To ensure that multipathing is configured correctly for your host, verify that you have the NetApp recommended settings configured for your ONTAP LUNs.

1. The AIX Host Utilities load the following parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Environment	Value for AIX	Note
algorithm	MPIO	round_robin	Set by Host Utilities
hcheck_cmd	MPIO	inquiry	Set by Host Utilities
hcheck_interval	MPIO	30	Set by Host Utilities
hcheck_mode	MPIO	nonactive	Set by Host Utilities
lun_reset_spt	MPIO / non-MPIO	yes	Set by Host Utilities
max_transfer	MPIO / non-MPIO	FC LUNs: 0x100000 bytes	Set by Host Utilities
qfull_dly	MPIO / non-MPIO	2-second delay	Set by Host Utilities
queue_depth	MPIO / non-MPIO	64	Set by Host Utilities
reserve_policy	MPIO / non-MPIO	no_reserve	Set by Host Utilities
re_timeout (disk)	MPIO / non-MPIO	30 seconds	Uses OS Default values
dyntrk	MPIO / non-MPIO	Yes	Uses OS Default values
fc_err_recov	MPIO / non-MPIO	Fast_fail	Uses OS Default values
q_type	MPIO / non-MPIO	simple	Uses OS Default values
num_cmd_elems	MPIO / non-MPIO	1024 for AIX 3072 for VIOS	FC EN1B, FC EN1C
num_cmd_elems	MPIO / non-MPIO	1024 for AIX	FC EN0G

2. If your storage configuration includes MetroCluster or SnapMirror active sync, change the default settings:

MetroCluster

By default, the AIX OS enforces a shorter I/O timeout when there are no available paths to a LUN. This might occur in configurations that include single-switch SAN fabric and in MetroCluster configurations that experience unplanned failovers. For additional information and recommended changes to default settings, see the Knowledge Base article What are AIX Host support considerations in a MetroCluster configuration?.

SnapMirror active sync

Beginning with ONTAP 9.11.1, SnapMirror active sync is supported for an AIX host. The primary cluster in an AIX configuration is the "active" cluster.

In an AIX configuration, failovers are disruptive. With each failover, you need to perform a re-scan on the host for I/O operations to resume.

Refer to the Knowledge Base article How to configure an AIX host for SnapMirror active sync.

3. Verify the path status for ONTAP LUNs:

```
sanlun lun show
```

The following example outputs show the correct path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configurations

An ASA configuration optimizes all paths to a given LUN, keeping them active ("primary"). This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# sanlun lun show -p |grep -p hdisk78
                    ONTAP Path:
vs_aix_clus:/vol/chataix_205p2_vol_en_1_7/jfs_205p2_lun_en
                           LUN: 37
                      LUN Size: 15g
                   Host Device: hdisk78
                          Mode: C
            Multipath Provider: AIX Native
        Multipathing Algorithm: round_robin
------ ------- ------ ------- --------- ----------
host   vserver  AIX                      AIX MPIO
path   path     MPIO   host    vserver     path
state  type     path   adapter LIF       priority
------ ------- ------ ------- --------- ----------
up     primary  path0  fcs0    fc_aix_1     1
up     primary  path1  fcs0    fc_aix_2     1
up     primary  path2  fcs1    fc_aix_3     1
up     primary  path3  fcs1    fc_aix_4     1
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the correct output for an ONTAP LUN with two Active/Optimized ("primary") paths and two Active/Non-Optimized ("secondary") paths:

Show example

```
# sanlun lun show -p |grep -p hdisk78
                    ONTAP Path:
vs_aix_clus:/vol/chataix_205p2_vol_en_1_7/jfs_205p2_lun_en
                           LUN: 37
                      LUN Size: 15g
                   Host Device: hdisk78
                          Mode: C
            Multipath Provider: AIX Native
        Multipathing Algorithm: round_robin
------- ---------- ------ ------- ---------- ----------
host    vserver    AIX                        AIX MPIO
path    path       MPIO   host    vserver         path
state   type       path   adapter LIF         priority
------- ---------- ------ ------- ---------- ----------
up      secondary  path0  fcs0    fc_aix_1        1
up      primary    path1  fcs0    fc_aix_2        1
up      primary    path2  fcs1    fc_aix_3        1
up      secondary  path3  fcs1    fc_aix_4        1
```

Step 4: Review the known issues

Known issues

The IBM AIX 7.2 and/or PowerVM (VIOS 3.1) with ONTAP storage release has the following known issues:

NetApp Bug ID	Title	Description	Partner ID
1416221	AIX 7200-05-01 encountered I/O disruption on virtual iSCSI disks(VIOS 3.1.1.x) during storage failover	I/O disruption can happen during storage failover operations on AIX 7.2 TL5 hosts on the virtual iSCSI disks mapped through the VIOS 3.1.1.x. By default, the `rw_timeout` value of the virtual iSCSI disks (hdisk) on VIOC will be 45 seconds. If an I/O delay greater than 45 seconds happens during storage failover, an I/O failure might occur. To avoid this situation, refer to the workaround mentioned in the BURT. As per IBM, after applying APAR - IJ34739 (upcoming release) we can dynamically change the rw_timeout value using the `chdev` command.	NA
1414700	AIX 7.2 TL04 encountered I/O disruption on virtual iSCSI disks(VIOS 3.1.1.x) during storage failover	I/O disruption can happen during storage failover operations on AIX 7.2 TL4 hosts on the virtual iSCSI disks mapped through the VIOS 3.1.1.x. By default, the `rw_timeout` value of vSCSI adapter on VIOC is 45 seconds. If an I/O delay of more than 45 seconds happens during a storage failover, I/O failure might occur. To avoid this situation, refer to the workaround mentioned in the BURT.	NA
1307653	I/O issues occur on VIOS 3.1.1.10 during SFO faults and straight I/O	On VIOS 3.1.1, I/O failures might occur on NPIV client disks that are backed by 16 GB or 32 GB FC adapters. Additionally, the `vfchost` driver might stop processing I/O requests from the client. Applying IBM APAR IJ22290 IBM APAR IJ23222 fixes the issue.	NA

What's next?

Learn about using the AIX Host Utilities tool.

Configure AIX 7.1 for FCP and iSCSI with ONTAP storage

The AIX Host Utilities software provides management and diagnostic tools for AIX hosts that are connected to ONTAP storage. When you install the AIX Host Utilities on an AIX 7.1 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability. If your configuration doesn't support SAN booting, you can use a local boot.

SAN boot

SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for an AIX/PowerVM host host. You can set up a SAN boot LUN to work in a AIX Multipath I/O (MPIO) environment that is using the FC protocol and running AIX Host Utilities with either the FC or FCoE protocol. The method you use to create a SAN boot LUN and install a new OS image in an AIX MPIO environment depends on the protocol that you are using.

Steps

1. Use the Interoperability Matrix Tool to verify that your AIX OS, protocol, and ONTAP version support SAN booting.
2. Follow the best practices for setting up a SAN boot in the vendor documentation.

Local boot

Perform a local boot by installing the AIX OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install the AIX Host Utilities

NetApp strongly recommends installing the AIX Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data. The MPIO package from the Host Utilities provides MPIO support for AIX and VIOS.

Installing the AIX Host Utilities provides additional timeout settings on your AIX host.

Install AIX Host Utilities 6.1.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with an AIX 7.1 host to manage ONTAP LUNs.

Multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues with the remaining paths. The AIX and PowerVM environments of the Host Utilities use the AIXs native multipathing solution (MPIO).

The Path Control Module (PCM) is responsible for controlling multiple paths for an AIX host. The PCM is a storage vendor supplied code that handles path management and is installed and enabled during the Host Utilities installation.

To ensure that multipathing is configured correctly for your host, verify that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. The AIX Host Utilities load the following parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Environment	Value for AIX	Note
algorithm	MPIO	round_robin	Set by Host Utilities
hcheck_cmd	MPIO	inquiry	Set by Host Utilities
hcheck_interval	MPIO	30	Set by Host Utilities
hcheck_mode	MPIO	nonactive	Set by Host Utilities
lun_reset_spt	MPIO / non-MPIO	yes	Set by Host Utilities
max_transfer	MPIO / non-MPIO	FC LUNs: 0x100000 bytes	Set by Host Utilities
qfull_dly	MPIO / non-MPIO	2-second delay	Set by Host Utilities
queue_depth	MPIO / non-MPIO	64	Set by Host Utilities
reserve_policy	MPIO / non-MPIO	no_reserve	Set by Host Utilities
re_timeout (disk)	MPIO / non-MPIO	30 seconds	Uses OS Default values
dyntrk	MPIO / non-MPIO	Yes	Uses OS Default values
fc_err_recov	MPIO / non-MPIO	Fast_fail	Uses OS Default values
q_type	MPIO / non-MPIO	simple	Uses OS Default values
num_cmd_elems	MPIO / non-MPIO	1024 for AIX	FC EN1B, FC EN1C
num_cmd_elems	MPIO / non-MPIO	500 for AIX (standalone/physical) 200 for VIOC	FC EN0G

2. If your storage configuration includes MetroCluster or SnapMirror active sync, change the default settings:

MetroCluster

By default, the AIX OS enforces a shorter I/O timeout when there are no available paths to a LUN. This might occur in configurations that include single-switch SAN fabric and in MetroCluster configurations that experience unplanned failovers. For additional information and recommended changes to default settings, see the Knowledge Base article What are AIX Host support considerations in a MetroCluster configuration?.

SnapMirror active sync

Beginning with ONTAP 9.11.1, SnapMirror active sync is supported for an AIX host. The primary cluster in an AIX configuration is the "active" cluster.

In an AIX configuration, failovers are disruptive. With each failover, you need to perform a re-scan on the host for I/O operations to resume.

Refer to the Knowledge Base article How to configure an AIX host for SnapMirror active sync.

3. Verify the path status for ONTAP LUNs:

```
sanlun lun show
```

The following example outputs show the correct path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configurations

An ASA configuration optimizes all paths to a given LUN, keeping them active ("primary"). This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# sanlun lun show -p |grep -p hdisk78
                    ONTAP Path:
vs_aix_clus:/vol/chataix_205p2_vol_en_1_7/jfs_205p2_lun_en
                           LUN: 37
                      LUN Size: 15g
                   Host Device: hdisk78
                          Mode: C
            Multipath Provider: AIX Native
        Multipathing Algorithm: round_robin
------ ------- ------ ------- --------- ----------
host   vserver  AIX                      AIX MPIO
path   path     MPIO   host    vserver     path
state  type     path   adapter LIF       priority
------ ------- ------ ------- --------- ----------
up     primary  path0  fcs0    fc_aix_1     1
up     primary  path1  fcs0    fc_aix_2     1
up     primary  path2  fcs1    fc_aix_3     1
up     primary  path3  fcs1    fc_aix_4     1
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the correct output for an ONTAP LUN with two Active/Optimized ("primary") paths and two Active/Non-Optimized ("secondary") paths:

Show example

```
# sanlun lun show -p |grep -p hdisk78
                    ONTAP Path:
vs_aix_clus:/vol/chataix_205p2_vol_en_1_7/jfs_205p2_lun_en
                           LUN: 37
                      LUN Size: 15g
                   Host Device: hdisk78
                          Mode: C
            Multipath Provider: AIX Native
        Multipathing Algorithm: round_robin
------- ---------- ------ ------- ---------- ----------
host    vserver    AIX                        AIX MPIO
path    path       MPIO   host    vserver         path
state   type       path   adapter LIF         priority
------- ---------- ------ ------- ---------- ----------
up      secondary  path0  fcs0    fc_aix_1         1
up      primary    path1  fcs0    fc_aix_2         1
up      primary    path2  fcs1    fc_aix_3         1
up      secondary  path3  fcs1    fc_aix_4         1
```

Step 4: Review the known issues

The AIX 7.1 with ONTAP storage release has no known issues.

What's next?

Learn about using the AIX Host Utilities tool.

CentOS

Configure CentOS 8.x for FCP and iSCSI for ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a CentOS 8.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with CentOS 8.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

   ```
   chkconfig multipathd on
   ```

   ```
   /etc/init.d/multipathd start
   ```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038303634722b4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:7:1    sdfi   130:64   active ready running
  |- 11:0:9:1    sdiy  8:288     active ready running
  |- 11:0:10:1  sdml 69:464   active ready running
  |- 11:0:11:1  sdpt  131:304  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a098038303634722b4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:6:35 sdwb  69:624  active ready running
| |- 16:0:5:35 sdun  66:752  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 15:0:0:35 sdaj  66:48   active ready running
  |- 15:0:1:35 sdbx  68:176  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

For CentOS Red Hat compatible kernel known issues, see the known issues for the Red Hat Enterprise Linux (RHEL) 8.x series.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about CentOS Linux Virtualization (KVM)

  CentOS Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Citrix

Configure Citrix Xenserver 8.4 for FCP and iSCSI with ONTAP storage

Configure Citrix Xenserver 8.4 for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

The Linux Host Utilities software package doesn't support Citrix Xenserver operating systems.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Confirm the multipath configuration for your host

You can use multipathing with Citrix Xenserver 8.4 to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exists. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the host operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the native Linux OS compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"3 queue_if_no_path pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
/sbin/mpathutil list
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
/usr/sbin/mpathutil status
show topology
create: 3600a098038315045572b5930646f4b63 dm-1 NETAPP  ,LUN C-
Mode
size=9.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 15:0:0:0  sdb  8:16    active ready running
  |- 15:0:1:0  sdc  8:32    active ready running
  |- 16:0:0:0  sdcf 69:48   active ready running
  `- 16:0:1:0  sdcg 69:64   active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
/usr/sbin/mpathutil status
show topology
create: 3600a098038315045572b5930646f4b63 dm-1 NETAPP  ,LUN C-
Mode
size=9.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
  |- 15:0:0:0  sdb  8:16    active ready running
  `- 15:0:1:0  sdc  8:32    active ready running
`-+- policy='service-time 0' prio=50 status=active
    |- 16:0:0:0  sdcf 69:48   active ready running
    `- 16:0:1:0  sdcg 69:64   active ready running
```

Step 3: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

   "sda" is the local SCSI disk that you want to add to the blacklist.

   An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 4: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 5: Review the known issues

There are no known issues.

Configure Citrix Hypervisor 8.2 for FCP and iSCSI with ONTAP storage

Configure Citrix Hypervisor 8.2 for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

The Linux Host Utilities software package doesn't support Citrix Hypervisor operating systems.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.
3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Confirm the multipath configuration for your host

You can use multipathing with Citrix Hypervisor 8.2 to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exists. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the host operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the native Linux OS compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"3 queue_if_no_path pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
/sbin/mpathutil list
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
/usr/sbin/mpathutil status
show topology
create: 3600a098038315045572b5930646f4b63 dm-1 NETAPP  ,LUN C-
Mode
size=9.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 15:0:0:0  sdb  8:16    active ready running
  |- 15:0:1:0  sdc  8:32    active ready running
  |- 16:0:0:0  sdcf 69:48   active ready running
  `- 16:0:1:0  sdcg 69:64   active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
/usr/sbin/mpathutil status
show topology
create: 3600a098038315045572b5930646f4b63 dm-1 NETAPP  ,LUN C-
Mode
size=9.0G features='4 queue_if_no_path pg_init_retries 50
retain_attached_hw_handle' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
  |- 15:0:0:0  sdb  8:16    active ready running
  `- 15:0:1:0  sdc  8:32    active ready running
`-+- policy='service-time 0' prio=50 status=active
    |- 16:0:0:0  sdcf 69:48   active ready running
    `- 16:0:1:0  sdcg 69:64   active ready running
```

Step 3: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

   "sda" is the local SCSI disk that you want to add to the blacklist.

   An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 4: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 5: Review the known issues

There are no known issues.

ESXi

Use VMware vSphere 8.x with ONTAP

You can configure ONTAP SAN host settings for the VMware vSphere 8.x release with FC, FCoE, and iSCSI protocols.

Hypervisor SAN booting

Before you begin

If you decide to use SAN booting, it must be supported by your configuration. You can use the Interoperability Matrix Tool to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

Steps

1. Map the SAN boot LUN to the host.
2. Verify that multiple paths are available.

   Multiple paths become available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify that the boot was successful.

Multipathing

ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP) that manages the sub-plugins, Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs). By default, these SATP rules are available in ESXi.

For ONTAP storage, the "VMW_SATP_ALUA" plugin is used by default with "VMW_PSP_RR" as a path selection policy (PSP). You can run the following command to confirm the PSP:

```
`esxcli storage nmp satp rule list -s VMW_SATP_ALUA`
```

Example output:

```
Name           Device  Vendor    Model         Driver  Transport  Options
-------------------------------------------------------------------------
-------
VMW_SATP_ALUA          LSI       INF-01-00
reset_on_attempted_reserve  system
VMW_SATP_ALUA          NETAPP
reset_on_attempted_reserve  system

Rule Group  Claim Options  Default PSP  PSP Options     Description
----------------------------------------------------------------
tpgs_on        VMW_PSP_MRU                 NetApp E-Series arrays with
ALUA support
tpgs_on        VMW_PSP_RR                  NetApp arrays with ALUA
support
```

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with higher priorities are Active/Optimized. This means they are serviced by the controller where the aggregate is located. The paths with lower priorities are active but non-optimized because they are served from a different controller. The non-optimized paths are only used when optimized paths are not available.

Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths.

```
# esxcli storage nmp device list -d naa.600a098038314869372454524395855
```

Example output:

```
naa.600a0980383148693724545244395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba4:C0:T0:L11, vmhba3:C0:T0:L11
   Is USB: false
```

```
# esxcli storage nmp path list -d naa.600a0980383148693724545244395855
```

Example output:

```
fc.20000024ff7f4a51:21000024ff7f4a51-fc.2009d039ea3ab21f:2003d039ea3ab21f-
naa.600a0980383148693724545244395855
   Runtime Name: vmhba4:C0:T0:L11
   Device: naa.600a0980383148693724545244395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1001,
TPG_state=AO,RTP_id=4,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff7f4a50:21000024ff7f4a50-fc.2009d039ea3ab21f:2002d039ea3ab21f-
naa.600a0980383148693724545244395855
   Runtime Name: vmhba3:C0:T0:L11
   Device: naa.600a0980383148693724545244395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383148693724545244395855)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1001,
TPG_state=AO,RTP_id=3,RTP_health=UP}
```

```
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff7f4a51:21000024ff7f4a51-fc.2009d039ea3ab21f:2001d039ea3ab21f-
naa.600a098038314869372454524395855
   Runtime Name: vmhba4:C0:T3:L11
   Device: naa.600a098038314869372454524395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314869372454524395855)
   Group State: active unoptimized
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1000,
TPG_state=ANO,RTP_id=2,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff7f4a50:21000024ff7f4a50-fc.2009d039ea3ab21f:2000d039ea3ab21f-
naa.600a098038314869372454524395855
   Runtime Name: vmhba3:C0:T3:L11
   Device: naa.600a098038314869372454524395855
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314869372454524395855)
   Group State: active unoptimized
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1000,
TPG_state=ANO,RTP_id=1,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

Unresolved directive in hu_vsphere_8.adoc - include::_include/hu/reuse_hu_asa_configuration.adoc[]

```
esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453
```

Example output:

```
naa.600a098038314962485d543078486c7a
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1001,TPG_state=AO}{TPG_id=1000,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0; lastPathIndex=3:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba4:C0:T0:L14, vmhba4:C0:T1:L14, vmhba3:C0:T0:L14,
vmhba3:C0:T1:L14
   Is USB: false
```

```
# esxcli storage nmp path list -d naa.600a098038314962485d543078486c7a
```

Example output:

```
fc.200034800d756a75:210034800d756a75-fc.2018d039ea936319:2015d039ea936319-
naa.600a098038314962485d543078486c7a
   Runtime Name: vmhba4:C0:T0:L14
   Device: naa.600a098038314962485d543078486c7a
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1000,
TPG_state=AO,RTP_id=2,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.200034800d756a75:210034800d756a75-fc.2018d039ea936319:2017d039ea936319-
naa.600a098038314962485d543078486c7a
   Runtime Name: vmhba4:C0:T1:L14
   Device: naa.600a098038314962485d543078486c7a
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1001,
```

```
TPG_state=AO,RTP_id=4,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.200034800d756a74:210034800d756a74-fc.2018d039ea936319:2014d039ea936319-
naa.600a098038314962485d543078486c7a
   Runtime Name: vmhba3:C0:T0:L14
   Device: naa.600a098038314962485d543078486c7a
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1000,
TPG_state=AO,RTP_id=1,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.200034800d756a74:210034800d756a74-fc.2018d039ea936319:2016d039ea936319-
naa.600a098038314962485d543078486c7a
   Runtime Name: vmhba3:C0:T1:L14
   Device: naa.600a098038314962485d543078486c7a
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038314962485d543078486c7a)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config: {TPG_id=1001,
TPG_state=AO,RTP_id=3,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

vVol

Virtual Volumes (vVols) are a VMware object type that corresponds to a Virtual Machine (VM) disk, its snapshots, and fast clones.

ONTAP tools for VMware vSphere includes the VASA Provider for ONTAP, which provides the integration point for a VMware vCenter to leverage vVols based storage. When you deploy the ONTAP tools Open Virtualization Appliance (OVA), it is automatically registered with the vCenter server and enables the VASA Provider.

When you create a vVols datastore using the vCenter UI, it guides you to create FlexVols as backup storage for the datastore. vVols within vVols datastores are accessed by ESXi hosts using a protocol endpoint (PE). In SAN environments, one 4MB LUN is created on each FlexVol in the datastore for use as a PE. A SAN PE is an administrative logical unit (ALU). vVols are subsidiary logical units (SLUs).

Standard requirements and best practices for SAN environments apply when using vVols, including (but not limited to) the following:

- Create at least one SAN LIF on each node per SVM you intend to use. The best practice is to create at least two per node, but no more than necessary.

- Eliminate any single point of failure. Use multiple VMkernel network interfaces on different network subnets that use NIC teaming when multiple virtual switches are used, or use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput.
- Configure zoning, VLANs, or both as required for host connectivity.
- Verify that all required initiators are logged into the target LIFs on the desired SVM.

> You must deploy ONTAP tools for VMware vSphere to enable the VASA Provider. The VASA Provider will manage all of your iGroup settings for you, therefore there is no need to create or manage iGroups in a vVols environment.

NetApp does not recommend changing any vVols settings from default at this time.

Refer to the Interoperability Matrix Tool for specific versions of ONTAP tools, or legacy VASA Provider for your specific versions of vSphere and ONTAP.

For detailed information on provisioning and managing vVols, refer to the ONTAP tools for VMware vSphere documentation, VMware vSphere with ONTAP, and Virtual Volumes (vVols) with ONTAP tools 10.

Recommended settings

ATS locking

ATS locking is **mandatory** for VAAI compatible storage and upgraded VMFS5 and is required for proper interoperability and optimal VMFS shared storage I/O performance with ONTAP LUNs. Refer to VMware documentation for details on enabling ATS locking.

Settings	Default	ONTAP Recommended	Description
HardwareAcceleratedLocking	1	1	Helps enable the use of Atomic Test and Set (ATS) locking
Disk IOPs	1000	1	IOPS limit: The Round Robin PSP defaults to an IOPS limit of 1000. In this default case, a new path is used after 1000 I/O operations are issued.
Disk/QFullSampleSize	0	32	The count of QUEUE FULL or BUSY conditions it takes before ESXi starts throttling.

> Enable `Space-alloc` setting for all the LUNs mapped to VMware vSphere for `UNMAP` to work. For more details, refer to ONTAP Documentation.

Guest OS timeouts

You can manually configure the virtual machines with the recommended guest OS tunings. After tuning updates, you must reboot the guest for the updates to take effect.

GOS timeout values:

Guest OS Type	Timeouts
Linux variants	disk timeout = 60
Windows	disk timeout = 60
Solaris	disk timeout = 60 busy retry = 300 not ready retry = 300 reset retry = 30 max.throttle = 32 min.throttle = 8

Validate the vSphere tunable

You can use the following command to verify the `HardwareAcceleratedLocking` setting.

```
esxcli system settings advanced list --option /VMFS3/HardwareAcceleratedLocking
```

```
   Path: /VMFS3/HardwareAcceleratedLocking
   Type: integer
   Int Value: 1
   Default Int Value: 1
   Min Value: 0
   Max Value: 1
   String Value:
   Default String Value:
   Valid Characters:
   Description: Enable hardware accelerated VMFS locking (requires
compliant hardware). Please see http://kb.vmware.com/kb/2094604 before
disabling this option.
```

Validate the disk IOPs setting

You can use the following command to verify the IOPs setting.

```
esxcli storage nmp device list -d naa.600a098038304731783f506670553355
```

```
naa.600a0980383047317783f506670553355
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383047317783f506670553355)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config: {policy=rr,
iops=1,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba4:C0:T0:L82, vmhba3:C0:T0:L82
   Is USB: false
```

Validate the QFullSampleSize

You can use the following command to verify the QFullSampleSize.

```
esxcli system settings advanced list --option /Disk/QFullSampleSize
```

```
   Path: /Disk/QFullSampleSize
   Type: integer
   Int Value: 32
   Default Int Value: 0
   Min Value: 0
   Max Value: 64
   String Value:
   Default String Value:
   Valid Characters:
   Description: Default I/O samples to monitor for detecting non-transient
queue full condition. Should be nonzero to enable queue depth throttling.
Device specific QFull options will take precedence over this value if set.
```

Known issues

The VMware vSphere 8.x with ONTAP release has the following known issues:

NetApp Bug ID	Title	Description
1543660	I/O error occurs when Linux VMs using vNVMe adapters encounter a long all paths down (APD) window	Linux VMs running vSphere 8.x and later and using virtual NVMe (vNVME) adapters encounter an I/O error because the vNVMe retry operation is disabled by default. To avoid a disruption on Linux VMs running older kernels during an all paths down (APD) or a heavy I/O load, VMware has introduced a tunable "VSCSIDisableNvmeRetry" to disable the vNVMe retry operation.

Related information

- VMware vSphere with ONTAP
- VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
- NetApp ONTAP with NetApp SnapMirror active sync with VMware vSphere Metro Storage Cluster (vMSC)

Use VMware vSphere 7.x with ONTAP

You can use ONTAP SAN host configuration settings for the vSphere 7.x release with FC, FCoE and iSCSI protocols.

Hypervisor SAN Booting

Before you begin

If you decide to use SAN booting, it must be supported by your configuration. You can use the Interoperability Matrix Tool to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

Steps

1. Map the SAN boot LUN to the host.
2. Verify that multiple paths are available.

   Multiple paths become available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify that the boot was successful.

Multipathing

ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP) that manages the sub-plugins, Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs). These SATP rules are available by default in ESXi.

For ONTAP storage, the "VMW_SATP_ALUA" plugin is used by default with "VMW_PSP_RR" as a path selection policy (PSP). You can run the following command to confirm the PSP.

esxcli storage nmp satp rule list -s VMW_SATP_ALUA

```
Name           Device  Vendor    Model         Driver  Transport  Options
-------------  ------  --------  ------------  ------  ---------
--------------------------
VMW_SATP_ALUA          NETAPP
reset_on_attempted_reserve


Rule Group  Claim Options  Default PSP  PSP Options     Description
 ----------  -------------  -----------  --------------  -----------
system      tpgs_on        VMW_PSP_RR                   NetApp arrays with
ALUA support
```

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with higher priorities are Active/Optimized. This means they are serviced by the controller where the aggregate is located. The paths with lower priorities are active but non-optimized because they are served from a different controller. The non-optimized paths are only used when optimized paths are not available.

Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths.

esxcli storage nmp device list -d naa.600a098038313530772b4d673979372f

```
naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=AO}{TPG_id=1001,TPG_state=ANO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=1:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba3:C0:T3:L21, vmhba4:C0:T2:L21
   Is USB: false
```

esxcli storage nmp path list -d naa.600a098038313530772b4d673979372f

```

```

```
fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200b00a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba3:C0:T2:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active unoptimized
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=ANO,RTP_id=29,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200700a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba3:C0:T3:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=25,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200800a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba4:C0:T2:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=26,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200c00a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba4:C0:T3:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active unoptimized
```

```
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=ANO,RTP_id=30,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

Unresolved directive in hu_vsphere_7.adoc - include::_include/hu/reuse_hu_asa_configuration.adoc[]

esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453

```
naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1001,TPG_state=AO}{TPG_id=1000,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=2:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba4:C0:T0:L9, vmhba3:C0:T1:L9, vmhba3:C0:T0:L9,
vmhba4:C0:T1:L9
   Is USB: false
```

esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453

```
fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:204a00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba4:C0:T0:L9
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=6,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201d00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba3:C0:T1:L9
```

```
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=3,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201b00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba3:C0:T0:L9
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=1,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:201e00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba4:C0:T1:L9
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=4,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

vVol

Virtual Volumes (vVols) are a VMware object type that corresponds to a Virtual Machine (VM) disk, and its snapshots and fast-clones.

ONTAP tools for VMware vSphere includes the VASA Provider for ONTAP, which provides the integration point for a VMware vCenter to leverage vVols based storage. When you deploy the ONTAP tools OVA, it is automatically registered with the vCenter server and enables the VASA Provider.

When you create a vVols datastore using the vCenter UI, it guides you to create FlexVols as backup storage for the datastore. vVols within a vVols datastores are access by ESXi hosts using a protocol endpoint (PE). In SAN environments, one 4MB LUN is created on each FlexVol in the datastore for use as a PE. A SAN PE is an

administrative logical unit (ALU). vVols are subsidiary logical units (SLUs).

Standard requirements and best practices for SAN environments apply when using vVols, including (but not limited to) the following:

1. Create at least one SAN LIF on each node per SVM you intend to use. The best practice is to create at least two per node, but no more than necessary.
2. Eliminate any single point of failure. use multiple VMkernel network interfaces on different network subnets that use NIC teaming when multiple virtual switches are used. Or use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput.
3. Configure zoning and/or VLANs as required for host connectivity.
4. Ensure all required initiators are logged into the target LIFs on the desired SVM.

You must deploy ONTAP tools for VMware vSphere to enabled the VASA Provider. The VASA Provider will manage all of your igroup settings for you, so there is no need to create or manage igroups in a vVols environment.

NetApp does not recommend changing any vVols settings from the default at this time.

Refer to the Interoperability Matrix Tool for specific versions of ONTAP tools, or legacy VASA Provider for your specific versions of vSphere and ONTAP.

For detailed information on provisioning and managing vVols, refer to the ONTAP tools for VMware vSphere documentation as well as VMware vSphere with ONTAP and Virtual Volumes (vVols) with ONTAP tools 10.

Recommended Settings

ATS Locking

ATS locking is **mandatory** for VAAI compatible storage and upgraded VMFS5, and is required for proper interoperability and optimal VMFS shared storage I/O performance with ONTAP LUNs. Refer to VMware documentation for details on enabling ATS locking.

Settings	Default	ONTAP Recommended	Description
HardwareAcceleratedLocking	1	1	Helps enable the use of Atomic Test and Set (ATS) locking
Disk IOPs	1000	1	IOPS limit: The Round Robin PSP defaults to an IOPS limit of 1000. In this default case, a new path is used after 1000 I/O operations are issued.
Disk/QFullSampleSize	0	32	The count of QUEUE FULL or BUSY conditions it takes before ESXi starts throttling.

Enable Space-alloc setting for all the LUN's mapped to VMware vSphere for UNMAP to work. For More details, refer to ONTAP Documentation.

Guest OS timeouts

You can manually configure the virtual machines with the recommended guest OS tunings. After tuning updates, you must reboot the guest for the updates to take effect.

GOS timeout values:

Guest OS Type	Timeouts
Linux variants	disk timeout = 60
Windows	disk timeout = 60
Solaris	disk timeout = 60 busy retry = 300 not ready retry = 300 reset retry = 30 max.throttle = 32 min.throttle = 8

Validating the vSphere tunable

Use the following command to verify the HardwareAcceleratedLocking setting.

```
esxcli system settings advanced list --option /VMFS3/HardwareAcceleratedLocking
```

```
   Path: /VMFS3/HardwareAcceleratedLocking
   Type: integer
   Int Value: 1
   Default Int Value: 1
   Min Value: 0
   Max Value: 1
   String Value:
   Default String Value:
   Valid Characters:
   Description: Enable hardware accelerated VMFS locking (requires
compliant hardware). Please see http://kb.vmware.com/kb/2094604 before
disabling this option.
```

Validate the disk IOPs setting

Use the following command to verify the IOPs setting.

```
esxcli storage nmp device list -d naa.600a098038304731783f506670553355
```

```
naa.600a0980383047317783f506670553355
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383047317783f506670553355)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba4:C0:T0:L82, vmhba3:C0:T0:L82
   Is USB: false
```

Validating the QFullSampleSize

Use the following command to verify the QFullSampleSize

`esxcli system settings advanced list --option /Disk/QFullSampleSize`

```
   Path: /Disk/QFullSampleSize
   Type: integer
   Int Value: 32
   Default Int Value: 0
   Min Value: 0
   Max Value: 64
   String Value:
   Default String Value:
   Valid Characters:
   Description: Default I/O samples to monitor for detecting non-transient
queue full condition. Should be nonzero to enable queue depth throttling.
Device specific QFull options will take precedence over this value if set.
```

Known issues

There are no known issues for the VMware vSphere 7.x with ONTAP release.

Related information

- VMware vSphere with ONTAP
- VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
- NetApp ONTAP with NetApp SnapMirror active sync with VMware vSphere Metro Storage Cluster (vMSC)

Use VMware vSphere 6.5 and 6.7 with ONTAP

You can use ONTAP SAN host configuration settings for the vSphere 6.5.x and 6.7.x releases with FC, FCoE and iSCSI protocols.

Hypervisor SAN Booting

Before you begin

If you decide to use SAN booting, it must be supported by your configuration. You can use the Interoperability Matrix Tool to verify that your OS, HBA, HBA firmware and the HBA boot BIOS, and ONTAP version are supported.

Steps

1. Map the SAN boot LUN to the host.
2. Verify that multiple paths are available.

   Multiple paths become available after the host OS is up and running on the paths.

3. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

4. Reboot the host to verify that the boot was successful.

Multipathing

ESXi provides an extensible multipathing module called Native Multipathing Plug-In (NMP) that manages the sub-plugins, Storage Array Type Plugins (SATPs), and Path Selection Plugins (PSPs). These SATP rules are available by default in ESXi.

For ONTAP storage, the "VMW_SATP_ALUA" plugin is used by default with "VMW_PSP_RR" as a path selection policy (PSP). You can run the following command to confirm the PSP:

`esxcli storage nmp satp rule list -s VMW_SATP_ALUA`

```
Name           Device  Vendor    Model     Driver  Transport  Options
-------------  ------  ------    -------   ------  ---------  -------
VMW_SATP_ALUA           LSI       INF-01-00
reset_on_attempted_reserve
VMW_SATP_ALUA           NETAPP
reset_on_attempted_reserve

Rule Group  Claim Options  Default PSP  PSP Options  Description
----------  -------------  -----------  -----------  ------------
system        tpgs_on         VMW_PSP_MRU       NetApp E-Series arrays
with ALUA support
system        tpgs_on          MW_PSP_RR          NetApp arrays with ALUA
support
```

Unresolved directive in hu_vsphere_6.adoc - include::_include/hu/reuse_hu_asa_configuration.adoc[]

esxcli storage nmp device list -d naa.600a098038304759563f4e7837574453

```
fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:204a00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba4:C0:T0:L9
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=6,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201d00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba3:C0:T1:L9
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=3,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d36:21000024ff171d36-fc.202300a098ea5e27:201b00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba3:C0:T0:L9
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=1,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000024ff171d37:21000024ff171d37-fc.202300a098ea5e27:201e00a098ea5e27-
naa.600a098038304759563f4e7837574453
   Runtime Name: vmhba4:C0:T1:L9
```

```
   Device: naa.600a098038304759563f4e7837574453
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038304759563f4e7837574453)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=AO,RTP_id=4,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

In the above example, LUN has been mapped from NetApp storage with 4 paths (4 active-optimized).

Non-ASA configurations

For non-ASA configurations, there should be two groups of paths with different priorities. The paths with higher priorities are Active/Optimized. This means they are serviced by the controller where the aggregate is located. The paths with lower priorities are active but non-optimized because they are served from a different controller. The non-optimized paths are only used when optimized paths are not available.

Example

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths.

`esxcli storage nmp path list -d naa.600a098038313530772b4d673979372f`

```
fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200b00a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba3:C0:T2:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active unoptimized
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=ANO,RTP_id=29,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8e:10000090fae0ec8e-fc.201000a098dfe3d1:200700a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba3:C0:T3:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
```

```
{TPG_id=1000,TPG_state=AO,RTP_id=25,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200800a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba4:C0:T2:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1000,TPG_state=AO,RTP_id=26,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.

fc.20000090fae0ec8f:10000090fae0ec8f-fc.201000a098dfe3d1:200c00a098dfe3d1-
naa.600a098038313530772b4d673979372f
   Runtime Name: vmhba4:C0:T3:L21
   Device: naa.600a098038313530772b4d673979372f
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a098038313530772b4d673979372f)
   Group State: active unoptimized
   Array Priority: 0
   Storage Array Type Path Config:
{TPG_id=1001,TPG_state=ANO,RTP_id=30,RTP_health=UP}
   Path Selection Policy Path Config: PSP VMW_PSP_RR does not support path
configuration.
```

In the above example, LUN has been mapped from NetApp storage with 4 paths (2 active-optimized and 2 active-unoptimized).

vVol

Virtual Volumes (vVols) are a VMware object type that corresponds to a Virtual Machine (VM) disk, and its snapshots and fast clones.

ONTAP tools for VMware vSphere includes the VASA Provider for ONTAP, which provides the integration point for a VMware vCenter to leverage vVols based storage. When you deploy the ONTAP tools OVA, it is automatically registered with the vCenter server and enables the VASA Provider.

When you create a vVols datastore using the vCenter UI, it guides you to create FlexVols as backup storage for the datastore. vVols within a vVols datastores are accessed by ESXi hosts using a protocol endpoint (PE). In SAN environments, one 4MB LUN is created on each FlexVol volume volume volume in the datastore for use as a PE. A SAN PE is an administrative logical unit (ALU); vVols are subsidiary logical units (SLUs).

Standard requirements and best practices for SAN environments apply when using vVols, including (but not limited to) the following:

1. Create at least one SAN LIF on each node per SVM you intend to use. The best practice is to create at least two per node, but no more than necessary.
2. Eliminate any single point of failure. Use multiple VMkernel network interfaces on different network subnets that use NIC teaming when multiple virtual switches are used or use multiple physical NICs connected to multiple physical switches to provide HA and increased throughput.
3. Configure zoning and/or VLANs as required for host connectivity.
4. Ensure all required initiators are logged into the target LIFs on the desired SVM.

> You must deploy ONTAP tools for VMware vSphere to enable the VASA Provider. The VASA Provider will manage all of your igroup settings for you, so there is no need to create or manage iGroups in a vVols environment.

NetApp doesn't recommend changing any vVols settings from default at this time.

Refer to the Interoperability Matrix Tool for specific versions of ONTAP tools, or legacy VASA Provider for your specific versions of vSphere and ONTAP.

For detailed information on provisioning and managing vVols, refer to the ONTAP tools for VMware vSphere documentation as well as VMware vSphere with ONTAP and Virtual Volumes (vVols) with ONTAP tools 10.

Recommended Settings

ATS Locking

ATS locking is **mandatory** for VAAI compatible storage and upgraded VMFS5 and is required for proper interoperability and optimal VMFS shared storage I/O performance with ONTAP LUNs. Refer to VMware documentation for details on enabling ATS locking.

Settings	Default	ONTAP Recommended	Description
HardwareAcceleratedLocking	1	1	Helps enable the use of Atomic Test and Set (ATS) locking
Disk IOPs	1000	1	IOPS limit: The Round Robin PSP defaults to an IOPS limit of 1000. In this default case, a new path is used after 1000 I/O operations are issued.
Disk/QFullSampleSize	0	32	The count of QUEUE FULL or BUSY conditions it takes before ESXi starts throttling.

> Enable Space-alloc setting for all the LUN's mapped to VMware vSphere for UNMAP to work. For more details, refer to ONTAP Documentation.

Guest OS timeouts

You can manually configure the virtual machines with the recommended guest OS tunings. After tuning updates, you must reboot the guest for the updates to take effect.

GOS timeout values:

Guest OS Type	Timeouts
Linux variants	disk timeout = 60
Windows	disk timeout = 60
Solaris	disk timeout = 60 busy retry = 300 not ready retry = 300 reset retry = 30 max.throttle = 32 min.throttle = 8

Validate the vSphere tunable

Use the following command to verify the `HardwareAcceleratedLocking` setting:

esxcli system settings advanced list --option /VMFS3/HardwareAcceleratedLocking

```
   Path: /VMFS3/HardwareAcceleratedLocking
   Type: integer
   Int Value: 1
   Default Int Value: 1
   Min Value: 0
   Max Value: 1
   String Value:
   Default String Value:
   Valid Characters:
   Description: Enable hardware accelerated VMFS locking (requires
compliant hardware). Please see http://kb.vmware.com/kb/2094604 before
disabling this option.
```

Validate the disk IOPs setting

Use the following command to verify the IOPs setting:

esxcli storage nmp device list -d naa.600a098038304731783f506670553355

```
naa.600a0980383047317783f506670553355
   Device Display Name: NETAPP Fibre Channel Disk
(naa.600a0980383047317783f506670553355)
   Storage Array Type: VMW_SATP_ALUA
   Storage Array Type Device Config: {implicit_support=on;
explicit_support=off; explicit_allow=on; alua_followover=on;
action_OnRetryErrors=off;
{TPG_id=1000,TPG_state=ANO}{TPG_id=1001,TPG_state=AO}}
   Path Selection Policy: VMW_PSP_RR
   Path Selection Policy Device Config:
{policy=rr,iops=1,bytes=10485760,useANO=0; lastPathIndex=0:
NumIOsPending=0,numBytesPending=0}
   Path Selection Policy Device Custom Config:
   Working Paths: vmhba4:C0:T0:L82, vmhba3:C0:T0:L82
   Is USB: false
```

Validate the QFullSampleSize

Use the following command to verify the QFullSampleSize:

`esxcli system settings advanced list --option /Disk/QFullSampleSize`

```
   Path: /Disk/QFullSampleSize
   Type: integer
   Int Value: 32
   Default Int Value: 0
   Min Value: 0
   Max Value: 64
   String Value:
   Default String Value:
   Valid Characters:
   Description: Default I/O samples to monitor for detecting non-transient
queue full condition. Should be nonzero to enable queue depth throttling.
Device specific QFull options will take precedence over this value if set.
```

Known issues

The VMware vSphere 6.5 and 6.7 with ONTAP release has the following known issues:

OS version	NetApp Bug ID	Title	Description
ESXi 6.5 and ESXi 6.7.x	1413424	WFC RDM luns fails during testing	Windows failover clustering raw device mapping between Windows Virtual Machines like Windows 2019, Windows 2016, and Windows 2012 across VMWare ESXi host failed during storage failover testing on all the 7-mode, C-cmode cluster controllers.
ESXi 6.5.x and ESXi 6.7.x	1256473	PLOGI issue seen during testing on Emulex adapters	

Related information

- VMware vSphere with ONTAP
- VMware vSphere 5.x, 6.x and 7.x support with NetApp MetroCluster (2031038)
- NetApp ONTAP with NetApp SnapMirror active sync with VMware vSphere Metro Storage Cluster (vMSC)

HP-UX

Configure HP-UX 11i v3 for FCP and iSCSI with ONTAP storage

The HP-UX Host Utilities software provides management and diagnostic tools for HP-UX hosts that are connected to ONTAP storage. When you install the HP-UX Host Utilities on a HP-UX 11i v3 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

Configure your host to use SAN booting to simplify deployment and improve scalability. SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for a HP-UX host. The Host Utilities support SAN booting with FC and FCoE protocols in HP-UX environments.

Steps

1. Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.
2. Follow the best practices for setting up a SAN boot in the HP-UX vendor documentation.

Step 2: Install the HP-UX Host Utilities

NetApp strongly recommends installing the HP-UX Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install HP-UX Host Utilities 6.0

Step 3: Confirm the multipath configuration for your host

Use multipathing with HP-UX 11i v3 to manage ONTAP LUNs. Multipathing allows you to configure multiple network paths between the host and storage system. If one path fails, traffic continues on the remaining paths.

After installing HP-UX Host Utilities, verify that you have the NetApp recommended settings configured for your ONTAP LUNs.

About this task

The HP-UX Host Utilities support Native Microsoft Multipath I/O (MPIO) and Veritas Dynamic Multipathing. The following steps are for the Native MPIO solution.

Steps

1. When you install the HP-UX Host Utilities, the following recommended default settings are automatically loaded for ONTAP LUNs.

   Show parameter settings

Parameter	Uses Default Value
transient_secs	120
leg_mpath_enable	TRUE
max_q_depth	8
path_fail_secs	120
load_bal_policy	Round_robin
lua_enabled	TRUE
esd_secs	30

2. Verify the parameter settings and path status for your ONTAP LUNs:

   ```
   sanlun lun show
   ```

   The default multipath parameters support ASA, AFF, and FAS configurations. The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# sanlun lun show -p vs39:/vol/hpux_vol_1_1/hpux_lun

                    ONTAP Path: vs39:/vol/hpux_vol_1_1/hpux_lun
                           LUN: 2
                      LUN Size: 30g
                   Host Device: /dev/rdisk/disk25
                          Mode: C
              Multipath Provider: None
-------- ---------- -------------------- ------- --------------
host     vserver    /dev/dsk
path     path       filename             host    vserver
state    type       or hardware path     adapter LIF
-------- ---------- -------------------- ------- --------------
up       primary    /dev/dsk/c4t0d2      fcd0    248_1c_hp
up       primary    /dev/dsk/c6t0d2      fcd0    246_1c_hp
up       primary    /dev/dsk/c10t0d2     fcd1    246_1d_hp
up       primary    /dev/dsk/c8t0d2      fcd1    248_1d_hp
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# sanlun lun show -p vs39:/vol/vol24_3_0/lun24_0
                    ONTAP Path: vs39:/vol/vol24_3_0/lun24_0
                           LUN: 37
                      LUN Size: 15g
                   Host Device: /dev/rdisk/disk942
                          Mode: C
              Multipath Policy: A/A
            Multipath Provider: Native
------- ---------- ------ ------- ---------- ----------
--------------
host    vserver     /dev/dsk                             HP A/A
path    path        filename          host    vserver   path
failover
state   type        or hardware path adapter    LIF     priority
------- ---------- ------ ------- ---------- ----------
--------------
up       primary    /dev/dsk/c39t4d5   fcd0    hpux_3     0
up       primary    /dev/dsk/c41t4d5   fcd1    hpux_4     0
up       secondary  /dev/dsk/c40t4d5   fcd0    hpux_3     1
up       secondary  /dev/dsk/c42t4d5   fcd1    hpux_4     1
```

Step 4: Review the known issues

The HP-UX 11i v3 with ONTAP storage release has the following known issues:

NetApp Bug ID	Title	Description	Partner ID
1447287	AUFO event on the isolated master cluster in SnapMirror active sync configuration causes temporary disruption on the HP-UX host	This issue occurs when there is an automatic unplanned failover (AUFO) event on the isolated master cluster in the SnapMirror active sync configuration. It might take more than 120 seconds for I/O to resume on the HP-UX host, but this might not cause any I/O disruption or error messages. This issue causes dual event failure because the connection between the primary and the secondary cluster is lost and the connection between the primary cluster and the mediator is also lost. This is considered a rare event, unlike other AUFO events.	NA
1344935	HP-UX 11.31 Host intermittently reporting path status incorrectly on ASA setup.	Path reporting issues with ASA configuration.	NA
1306354	HP-UX LVM creation sends I/O of block size above 1MB	SCSI Maximum Transfer Length of 1MB is enforced in ONTAP All SAN Array. To restrict the Maximum Transfer Length from HP-UX hosts when connected to ONTAP All SAN Array, it is required to set the Maximum I/O size allowed by the HP-UX SCSI subsystem to 1MB. Refer HP-UX vendor documentation for details.	NA

What's next

Learn about using the HP-UX Host Utilities tool.

HPE VME

Configure HPE VME 8.0.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a HPE VME 8.0.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 2: Confirm the multipath configuration for your host

You can use multipathing with HPE VME 8.0.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a098038314c4a433f577471797958 dm-2 NETAPP,LUN C-Mode
size=180G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 14:0:0:0  sdc  8:32   active ready running
  |- 17:0:0:0  sdas 66:192 active ready running
  |- 14:0:3:0  sdar 66:176 active ready running
  `- 17:0:3:0  sdch 69:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren't available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 3: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

   ```
   $apt list |grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
days ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
   Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
             ├─1127418 /usr/sbin/iscsid
             └─1127419 /usr/sbin/iscsid
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

Step 4: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 5: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> ⚠ You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 6: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Configure HPE VME data stores on the ONTAP LUNs, then configure the virtual machine. Refer to the vendor HPE documentation for more information.

Oracle Linux

Configure Oracle Linux 9.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on an Oracle Linux 9.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Oracle Linux 9.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038303634722b4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
  |- 11:0:7:6   sdbz 68:208  active ready running
  |- 11:0:11:6  sddn 71:80   active ready running
  |- 11:0:15:6  sdfb 129:208 active ready running
  |- 12:0:1:6   sdgp 132:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren't available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383036347ffb4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:6:35 sdwb  69:624  active ready running
| |- 16:0:5:35 sdun  66:752  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 15:0:0:35 sdaj  66:48   active ready running
  |- 15:0:1:35 sdbx  68:176  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

Oracle Linux 9.3, 9.2, 9.1, and 9.0 with ONTAP storage has the following known issue:

NetApp Bug ID	Title	Description
1508554	SAN LUN utility with Emulex HBA needs symbolic links from library packages	When you execute the Linux Host Utilities CLI command - "sanlun fcp show adapter -v" on a SAN host, the command fails with an error message displaying that the library dependencies required for an host bus adapter (HBA) discovery cannot be located: ---- [root@hostname ~]# sanlun fcp show adapter -v Unable to locate /usr/lib64/libHBAAPI.so library Make sure the package installing the library is installed & loaded ----

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Oracle Linux KVM and Virtualization

  Oracle Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. You can manage and support multiple Oracle Linux KVM hosts by using Oracle Linux Virtualization Manager, which is built from the open-source oVirt project. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure Oracle Linux 8.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on an Oracle Linux 8.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Oracle Linux 8.x to manage ONTAP LUNs.

You can use the recommended settings for Red Hat Enterprise Linux (RHEL) 8.x to configure Red Hat Compatible Kernel for Oracle Linux 8.x.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

   ```
   systemctl enable multipathd
   ```

   ```
   systemctl start multipathd
   ```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038303634722b4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
  |- 11:0:7:6   sdbz 68:208  active ready running
  |- 11:0:11:6  sddn 71:80   active ready running
  |- 11:0:15:6  sdfb 129:208 active ready running
  |- 12:0:1:6   sdgp 132:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383036347ffb4d59646c4436 dm-28 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:6:35 sdwb  69:624  active ready running
| |- 16:0:5:35 sdun  66:752  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 15:0:0:35 sdaj  66:48   active ready running
  |- 15:0:1:35 sdbx  68:176  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Oracle Linux KVM and Virtualization

  Oracle Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. You can manage and support multiple Oracle Linux KVM hosts by using Oracle Linux Virtualization Manager, which is built from the open-source oVirt project. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Proxmox

Configure Proxmox VE 9.x for FCP and iSCSI with ONTAP storage

Configure Proxmox VE 9.x for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

FCP and iSCSI with Proxmox VE 9.x have the following known limitations:

- The Linux Host Utilities do not support Proxmox VE 9.x operating systems.
- The SAN boot configuration is not supported.

Step 1: Confirm the multipath configuration for your host

You can use multipathing with Proxmox VE 9.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"always"
hardware_handler	"1"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Optionally, override the default value for the `find_multipaths` parameter to ensure that ONTAP LUNs are correctly discovered and managed by multipathd:

   a. Set `find_multipaths` to "no" in the defaults section of `/etc/multipath.conf`:

```
defaults {
    find_multipaths "no"
}
```

b. Reload the multipath service:

```
systemctl reload multipathd
```

> By default, the Proxmox OS-native multipath configuration sets `find_multipaths` to "**strict**" with the empty zero byte `/etc/multipath.conf` configuration file each time you boot the host. This can prevent the host discovering newly presented ONTAP LUNs as multipath devices, which means they do not appear under multipath control automatically. Existing ONTAP LUNs remain discovered and under multipath control after each reboot.

5. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. If there are more than four paths, it might cause issues with the paths during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038315071592b59713261566d dm-38 NETAPP,LUN C-Mode
size=100G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 8:0:0:7 sdbv 68:144 active ready running
  |- 9:0:0:7 sdbx 68:176 active ready running
  |- 6:0:0:7 sdbr 68:80  active ready running
  `- 7:0:0:7 sdbt 68:112 active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example shows the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 2: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

   ```
   $apt list |grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
days ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
   Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
             ├─1127418 /usr/sbin/iscsid
             └─1127419 /usr/sbin/iscsid
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

Step 3: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 4: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> ⚠ You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 5: Review the known issues

There are no known issues.

Configure Proxmox VE 8.x for FCP and iSCSI with ONTAP storage

Configure Proxmox VE 8.x for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

FCP and iSCSI with Proxmox VE 8.x have the following known limitations:

- The Linux Host Utilities do not support Proxmox VE 8.x operating systems.
- The SAN boot configuration is not supported.

Step 1: Confirm the multipath configuration for your host

You can use multipathing with Proxmox VE 8.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"always"
hardware_handler	"1"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Optionally, override the default value for the `find_multipaths` parameter to ensure that ONTAP LUNs are correctly discovered and managed by multipathd:

   a. Set `find_multipaths` to "no" in the defaults section of `/etc/multipath.conf`:

   ```
   defaults {
       find_multipaths "no"
   }
   ```

   b. Reload the multipath service:

   ```
   systemctl reload multipathd
   ```

By default, the Proxmox OS-native multipath configuration sets `find_multipaths` to "**strict**" with the empty zero byte `/etc/multipath.conf` configuration file each time you boot the host. This can prevent the host discovering newly presented ONTAP LUNs as multipath devices, which means they do not appear under multipath control automatically. Existing ONTAP LUNs remain discovered and under multipath control after each reboot.

5. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. If there are more than four paths, it might cause issues with the paths during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038315071592b59713261566d dm-38 NETAPP,LUN C-Mode
size=100G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 8:0:0:7 sdbv 68:144 active ready running
  |- 9:0:0:7 sdbx 68:176 active ready running
  |- 6:0:0:7 sdbr 68:80  active ready running
  `- 7:0:0:7 sdbt 68:112 active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example shows the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 2: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

   ```
   $apt list |grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
days ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
   Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
             ├─1127418 /usr/sbin/iscsid
             └─1127419 /usr/sbin/iscsid
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

Step 3: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 4: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> ⚠ You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 5: Review the known issues

There are no known issues.

RHEL

Configure RHEL 10.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a Red Hat Enterprise Linux (RHEL) 10.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with RHEL 10.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a098038314e535a24584e4b496252 dm-32 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:0:41 sdan 66:112  active ready running
  |- 11:0:1:41 sdcb 68:240  active ready running
  |- 14:0:2:41 sdfd 129:240 active ready running
  `- 14:0:0:41 sddp 71:112  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Red Hat Linux Virtualization (KVM)

  Red Hat Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure RHEL 9.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a Red Hat Enterprise Linux (RHEL) 9.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with RHEL 9.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038314c4a433f577471797958 dm-2 NETAPP,LUN C-Mode
size=180G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 14:0:0:0  sdc  8:32   active ready running
  |- 17:0:0:0  sdas 66:192 active ready running
  |- 14:0:3:0  sdar 66:176 active ready running
  `- 17:0:3:0  sdch 69:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

RHEL 9.x with ONTAP storage has the following known issues.

9.3

NetApp Bug ID	Title	Description	JIRA ID
1508554	NetApp Linux Host Utilities CLI requires additional library package dependencies to support Emulex host bus adapter (HBA) adapter discovery	In RHEL 9.x, the NetApp Linux SAN Host Utilities CLI `sanlun fcp show adapter -v` fails because the library package dependencies to support Emulex host bus adapter (HBA) discovery cannot be found.	Not applicable
1593771	A Red Hat Enterprise Linux 9.3 QLogic SAN host encounters loss of partial multipaths during storage mobility operations	During the ONTAP storage controller takeover operation, half of the multipaths are expected to go down or switch to a failover mode and then recover to full path count during the giveback workflow. However, with a Red Hat Enterprise Linux (RHEL) 9.3 QLogic host, only partial multipaths are recovered after a storage failover giveback operation.	RHEL 17811

9.2

NetApp Bug ID	Title	Description
1508554	NetApp Linux Host Utilities CLI requires additional library package dependencies to support Emulex HBA adapter discovery	In RHEL 9.2, the NetApp Linux SAN Host Utilities CLI `sanlun fcp show adapter -v` fails because the library package dependencies to support HBA discovery cannot be found.
1537359	A Red Hat Linux 9.2 SAN booted host with Emulex HBA encounters stalled tasks leading to kernel disruption	During a storage failover giveback operation, a Red Hat Linux 9.2 SAN booted host with an Emulex host bus adapeter (HBA) encounters stalled tasks leading to kernel disruption. The kernel disruption causes the operating system to reboot and if `kdump` is configured, it generates the `vmcore` file under the `/var/crash/` directory. The issue is being triaged with the `lpfc` driver but it cannot be reproduced consistently.

9.1

NetApp Bug ID	Title	Description
1508554	NetApp Linux Host Utilities CLI requires additional library package dependencies to support Emulex HBA adapter discovery	In RHEL 9.1, the NetApp Linux SAN Host Utilities CLI `sanlun fcp show adapter -v` fails because the library package dependencies to support HBA discovery cannot be found.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Red Hat Linux Virtualization (KVM)

  Red Hat Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure RHEL 8.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a Red Hat Enterprise Linux (RHEL) 8.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with RHEL 8.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038314c4a433f577471797958 dm-2 NETAPP,LUN C-Mode
size=180G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 14:0:0:0  sdc  8:32   active ready running
  |- 17:0:0:0  sdas 66:192 active ready running
  |- 14:0:3:0  sdar 66:176 active ready running
  `- 17:0:3:0  sdch 69:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

RHEL 8.x with ONTAP storage has the following known issues.

8.1

NetApp Bug ID	Title	Description
1275843	Kernel disruption might occur on Red Hat Enterprise Linux 8.1 with QLogic QLE2672 16GB FC HBA during storage failover operation	Kernel disruption might occur during storage failover operations on the Red Hat Enterprise Linux 8.1 kernel with a QLogic QLE2672 Fibre Channel (FC) host bus adapter (HBA). The kernel disruption causes Red Hat Enterprise Linux 8.1 to reboot, leading to application disruption. If the kdump mechanism is enabled, the kernel disruption generates a vmcore file located in the/var/crash/ directory. You can check the vmcore file to determine the cause of the disruption.A storage failover with the QLogic QLE2672 HBA event affects the "kmem_cache_alloc+131" module. You can locate the event in the vmcore file by finding the following string: " [exception RIP: kmem_cache_alloc+131]"After the kernel disruption, reboot the Host OS and recover the operating system. Then restart the applications

NetApp Bug ID	Title	Description
1275838	Kernel disruption occurs on Red Hat Enterprise Linux 8.1 with QLogic QLE2742 32GB FC HBA during storage failover operations	Kernel disruption occurs during storage failover operations on the Red Hat Enterprise Linux 8.1 kernel with a QLogic QLE2742 Fibre Channel (FC) host bus adapter (HBA). The kernel disruption causes Red Hat Enterprise Linux 8.1 to reboot, leading to application disruption. If the kdump mechanism is enabled, the kernel disruption generates a vmcore file located in the/var/crash/ directory. You can check the vmcore file to determine the cause of the disruption. A storage failover with the QLogic QLE2742 HBA event affects the "kmem_cache_alloc+131"module. You can locate the event in the vmcore file by finding the following string: " [exception RIP: kmem_cache_alloc+131]"After the kernel disruption, reboot the Host OS and recover the operating system. Then restart the applications.
1266250	Login to multiple paths fails during the Red Hat Enterprise Linux 8.1 installation on iSCSI SAN LUN	You cannot login to multiple paths during the Red Hat Enterprise Linux 8.1 installation on iSCSI SAN LUN multipath devices. Installation is not possible on multipath iSCSI device and the multipath service is not enabled on the SAN boot device.

8.0

NetApp Bug ID	Title	Description
1238719	Kernel disruption on RHEL8 with QLogic QLE2672 16GB FC during storage failover operations	Kernel disruption might occur during storage failover operations on a Red Hat Enterprise Linux (RHEL) 8 kernel with a QLogic QLE2672 host bus adapter (HBA). The kernel disruption causes the operating system to reboot. The reboot causes application disruption and generates the vmcore file under the /var/crash/directory if kdump is configured. Use the vmcore file to identify the cause of the failure. In this case, the disruption is in the “kmem_cache_alloc+160” module. It is logged in the vmcore file with the following string: "[exception RIP: kmem_cache_alloc+160]". Reboot the host OS to recover the operating system and then restart the application.
1226783	RHEL8 OS boots up to "emergency mode" when more than 204 SCSI devices are mapped on all Fibre Channel (FC) host bus adapters (HBA)	If a host is mapped with more than 204 SCSI devices during an operating systemreboot process, the RHEL8 OS fails to boot up to "normal mode" and enters "emergency mode". This results in most of the host services becoming unavailable.
1230882	Creating a partition on an iSCSI multipath device during the RHEL8 installation is not feasible.	iSCSI SAN LUN multipath devices are not listed in disk selection during RHEL 8 installation. Consequently, the multipath service is not enabled on the SAN boot device.
1235998	The "rescan-scsi-bus.sh -a" command does not scan more than 328 devices	If a Red Hat Enterprise Linux 8 host maps with more than 328 SCSI devices, the host OS command "rescan-scsi-bus.sh -a" only scans 328 devices. The host does not discover any remaining mapped devices.

NetApp Bug ID	Title	Description
1231087	Remote ports transit to a blocked state on RHEL8 with Emulex LPe16002 16GB FC during storage failover operations	Remote ports transit to a blocked state on RHEL8 with Emulex LPe16002 16GB Fibre Channel (FC) during storage failover operations. When the storage node returns to an optimal state, the LIFs also come up and the remote port state should read "online". Occasionally, the remote port state might continue to read as "blocked" or "not present". This state can lead to a "failed faulty" path to LUNs at the multipath layer
1231098	Remote ports transit to blocked state on RHEL8 with Emulex LPe32002 32GB FC during storage failover operations	Remote ports transit to a blocked state on RHEL8 with Emulex LPe32002 32GBFibre Channel (FC) during storage failover operations. When the storage node returns to an optimal state, the LIFs also come up and the remote port state should read "online". Occasionally, the remote port state might continue to read as "blocked" or "not present". This state can lead to a "failed faulty" path to LUNs at the multipath layer.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Red Hat Linux Virtualization (KVM)

  Red Hat Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Rocky Linux

Configure Rocky Linux 10.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install Linux Host Utilities on a Rocky Linux 10.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Rocky Linux 10.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a098038314e535a24584e4b496252 dm-32 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:0:41 sdan 66:112  active ready running
  |- 11:0:1:41 sdcb 68:240  active ready running
  |- 14:0:2:41 sdfd 129:240 active ready running
  `- 14:0:0:41 sddp 71:112  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Rocky Linux Virtualization (KVM)

  Rocky Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure Rocky Linux 9.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install Linux Host Utilities on a Rocky Linux 9.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Rocky Linux 9.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
multipath -ll
3600a098038314c4a433f577471797958 dm-2 NETAPP,LUN C-Mode
size=180G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 14:0:0:0  sdc  8:32   active ready running
  |- 17:0:0:0  sdas 66:192 active ready running
  |- 14:0:3:0  sdar 66:176 active ready running
  `- 17:0:3:0  sdch 69:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren't available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

Rocky Linux 9.x with ONTAP storage has the following known issues.

9.3

NetApp Bug ID	Title	Description	JIRA ID
1508554	NetApp Linux Host Utilities CLI requires additional library package dependencies to support Emulex host bus adapter (HBA) adapter discovery	In RHEL 9.x, the NetApp Linux SAN Host Utilities CLI `sanlun fcp show adapter -v` fails because the library package dependencies to support Emulex host bus adapter (HBA) discovery cannot be found.	Not applicable
1593771	A Red Hat Enterprise Linux 9.3 QLogic SAN host encounters loss of partial multipaths during storage mobility operations	During the ONTAP storage controller takeover operation, half of the multipaths are expected to go down or switch to a failover mode and then recover to full path count during the giveback workflow. However, with a Red Hat Enterprise Linux (RHEL) 9.3 QLogic host, only partial multipaths are recovered after a storage failover giveback operation.	RHEL 17811

9.2

NetApp Bug ID	Title	Description
1508554	NetApp Linux Host Utilities CLI requires additional library package dependencies to support Emulex HBA adapter discovery	In RHEL 9.2, the NetApp Linux SAN Host Utilities CLI `sanlun fcp show adapter -v` fails because the library package dependencies to support HBA discovery cannot be found.
1537359	A Red Hat Linux 9.2 SAN booted host with Emulex HBA encounters stalled tasks leading to kernel disruption	During a storage failover giveback operation, a Red Hat Linux 9.2 SAN booted host with an Emulex host bus adapeter (HBA) encounters stalled tasks leading to kernel disruption. The kernel disruption causes the operating system to reboot and if `kdump` is configured, it generates the `vmcore` file under the `/var/crash/` directory. The issue is being triaged with the `lpfc` driver but it cannot be reproduced consistently.

9.1

NetApp Bug ID	Title	Description
1508554	NetApp Linux Host Utilities CLI requires additional library package dependencies to support Emulex HBA adapter discovery	In Rocky Linux 9.1, the NetApp Linux SAN Host Utilities CLI `sanlun fcp show adapter -v` fails because the library package dependencies to support HBA discovery cannot be found.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Rocky Linux Virtualization (KVM)

  Rocky Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure Rocky Linux 8.x for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install Linux Host Utilities on a Rocky Linux 8.x host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Rocky Linux 8.x to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exists:

   ```
   ls /etc/multipath.conf
   ```

   If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

   ```
   systemctl enable multipathd
   ```

   ```
   systemctl start multipathd
   ```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the host operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the native Linux OS compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a098038314c4a433f577471797958 dm-2 NETAPP,LUN C-Mode
size=180G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 14:0:0:0  sdc  8:32   active ready running
  |- 17:0:0:0  sdas 66:192 active ready running
  |- 14:0:3:0  sdar 66:176 active ready running
  `- 17:0:3:0  sdch 69:80  active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren't available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383149764b5d567257516273 dm-0 NETAPP,LUN C-Mode
size=150G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 16:0:3:0  sdcg 69:64  active ready running
| `- 10:0:0:0  sdb  8:16   active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 10:0:1:0  sdc  8:32   active ready running
  `- 16:0:2:0  sdcf 69:48  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep iscsi-initiator-utils
   ```

   You should see an output similar to the following example:

   ```
   iscsi-initiator-utils-6.2.1.11-0.git4b3e853.el9.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Rocky Linux Virtualization (KVM)

  Rocky Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Solaris

Configure Solaris 11.4 for FCP and iSCSI with ONTAP storage

The Solaris Host Utilities software provides management and diagnostic tools for Solaris hosts that are connected to ONTAP storage. When you install the Solaris Host Utilities on a Solaris 11.4 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability. If your configuration doesn't support SAN booting, you can use a local boot.

SAN boot

SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for a Solaris host. You can set up a SAN boot LUN to work in a Solaris MPxIO environment that is using the FC protocol and running Solaris Host Utilities. The method you use to set up a SAN boot LUN depends on your volume manager and file system.

Steps

1. Use the Interoperability Matrix Tool to verify that your Solaris OS, protocol, and ONTAP version support SAN booting.
2. Follow the best practices for setting up a SAN boot in the Solaris vendor documentation.

Local boot

Perform a local boot by installing the Solaris OS on the local hard disk, for example, install on an SSD, SATA, or RAID.

Step 2: Install the Solaris Host Utilities

NetApp strongly recommends installing the Solaris Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Installing the Solaris Host Utilities changes some of the timeout settings on your Solaris host.

Install Solaris Host Utilities 8.0.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Solaris 11.4 to manage ONTAP LUNs.

Multipathing enables you to configure multiple network paths between the host and storage systems. If one path fails, traffic continues on the remaining paths. Oracle Solaris I/O Multipathing (MPxIO) is enabled by default for Solaris 11.4 and SPARC systems.

Steps

1. If your host is configured for FC, verify that the default setting in `/kernel/drv/fp.conf` is set to `mpxio-disable="no"`.
2. The Solaris Host Utilities load the NetApp recommended parameter settings for SPARC and x86_64 processors.

   Show parameter settings

Parameter	Value
throttle_max	8
not_ready_retries	300
busy_retries	30
reset_retries	30
throttle_min	2
timeout_retries	10
physical_block_size	4096
disksort	false
cache-nonvolatile	true

   For additional information about Solaris 11.4 system settings, see Oracle Support DOC ID: 2595926.1.

3. If your storage configuration includes MetroCluster, Oracle Solaris virtualization, or SnapMirror active sync, change the default settings:

MetroCluster

By default, the Solaris OS fails to execute the I/O operations after **20s** if all paths to a LUN are lost. This is controlled by the `fcp_offline_delay` parameter. The default value for `fcp_offline_delay` is the appropriate for standard ONTAP clusters. However, in MetroCluster configurations, you must increase the value of `fcp_offline_delay` to **120s** to ensure that I/O doesn't prematurely time out during operations, including unplanned failovers.

For additional information and recommended changes to default settings for MetroCluster, see the Knowledge Base article Solaris host support considerations in a MetroCluster configuration.

Oracle Solaris virtualization

- Solaris virtualization options include Solaris Logical Domains (also called LDOMs or Oracle VM Server for SPARC), Solaris Dynamic Domains, Solaris Zones, and Solaris Containers. These technologies are also referred to as "Oracle Virtual Machines".
- You can use multiple options together, for example, a Solaris Container within a particular Solaris Logical Domain.
- NetApp supports the use of Solaris virtualization technologies where the overall configuration is supported by Oracle and any partition with direct access to LUNs is listed on the IMT in a supported configuration. This includes root containers, LDOM I/O domains, and LDOM using NPIV to access LUNs.
- Partitions or virtual machines that use only virtualized storage resources, such as a `vdsk`, don't need specific qualifications because they don't have direct access to ONTAP LUNs. You only need to verify that the partition or virtual machine that has direct access to the underlying LUN, such as an LDOM I/O domain, is listed on the IMT.

Steps

When LUNs are used as virtual disk devices within an LDOM, the source of the LUN is masked by virtualization and the LDOM doesn't correctly detect the block sizes. To prevent this issue:

a. Patch the LDOM OS for *Oracle Bug 15824910*

b. Create a `vdc.conf` file that sets the block size of the virtual disk to `4096`. See Oracle DOC: 2157669.1 for more information.

c. Verify the patch installation to ensure that the recommended settings have been configured correctly:

   i. Create a zpool:

   ```
   zpool create zpool_name disk_list
   ```

   ii. Run `zdb -C` against the zpool and verify that the value of **ashift** is `12`.

   If the value of **ashift** isn't `12`, re run `zdb -C11`, and verify that the correct patch was installed and recheck the contents of `vdc.conf`.

   Don't proceed until **ashift** shows a value of `12`.

Patches are available for Oracle bug 15824910 on several Solaris versions. Contact Oracle if you need assistance in determining the best kernel patch.

SnapMirror active sync

Beginning with ONTAP 9.9.1, SnapMirror active sync setting configurations are supported in the Solaris host. To verify that the Solaris client applications are non-disruptive when an unplanned site failover switchover occurs in a SnapMirror active sync environment, you must configure the `scsi-vhci-failover-override` setting on the Solaris host. This setting overrides the failover module `f_tpgs` to prevent the execution of the code path that detects the contradiction.

Steps

a. Create the configuration file `/etc/driver/drv/scsi_vhci.conf` with an entry similar to the following example for the NetApp storage type connected to the host:

```
scsi-vhci-failover-override =
"NETAPP  LUN","f_tpgs"
```

b. Verify that the override parameter has been successfully applied:

```
devprop
```

```
mdb
```

Show examples

```
root@host-A:~# devprop -v -n /scsi_vhci scsi-vhci-failover-
override      scsi-vhci-failover-override=NETAPP  LUN + f_tpgs
root@host-A:~# echo "*scsi_vhci_dip::print -x struct dev_info
devi_child | ::list struct dev_info devi_sibling| ::print
struct dev_info devi_mdi_client| ::print mdi_client_t
ct_vprivate| ::print struct scsi_vhci_lun svl_lun_wwn
svl_fops_name"| mdb -k
```

```
svl_lun_wwn = 0xa002a1c8960 "600a098038313477543f524539787938"
svl_fops_name = 0xa00298d69e0 "conf f_tpgs"
```

After `scsi-vhci-failover-override` has been applied, `conf` is added to `svl_fops_name`. For additional information and recommended changes to default settings, refer to the NetApp Knowledge Base article Solaris Host support recommended settings in SnapMirror active sync configuration.

4. Verify that 4KB aligned I/O with zpools using ONTAP LUNs is supported:

a. Verify that your Solaris host is installed with the latest Support Repository Update (SRU):

```
pkg info entire`
```

b. Verify that the ONTAP LUN has `ostype` as "Solaris", independent of the LUN size:

```
lun show -vserver` <vsersver_name>
```

Show example

```
chat-a800-31-33-35-37::*> lun show -vserver solaris_fcp -path
/vol/sol_195_zpool_vol_9/lun -fields ostype
vserver     path                          ostype
----------- ----------------------------- -------
solaris_fcp /vol/sol_195_zpool_vol_9/lun solaris
```

5. Verify the output for your ONTAP LUNs:

```
sanlun lun show
```

You should see an output similar to the following example for an ASA, AFF, or FAS configuration:

Show example

```
root@sparc-s7-55-148:~# sanlun lun show -pv

                    ONTAP Path: Solaris_148_siteA:/vol/Triage/lun
                           LUN: 0
                      LUN Size: 20g
                   Host Device:
/dev/rdsk/c0t600A098038314B32685D573064776172d0s2
                          Mode: C
              Multipath Provider: Sun Microsystems
                Multipath Policy: Native
```

6. Verify the path status for your ONTAP LUNs:

```
mpathadm show lu <LUN>`
```

The following example outputs show the correct path status for ONTAP LUNs in an ASA, AFF, or FAS configuration. The path priorities are displayed against "Access State" for each LUN in the output.

ASA configurations

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
root@sparc-s7-55-82:~# mpathadm show lu
/dev/rdsk/c0t600A098038313953495D58674777794Bd0s2
Logical Unit:  /dev/rdsk/c0t600A098038313953495D58674777794Bd0s2
        mpath-support:  libmpscsi_vhci.so
        Vendor:  NETAPP
        Product:  LUN C-Mode
        Revision:  9171
        Name Type:  unknown type
        Name:  600a098038313953495d58674777794b
        Asymmetric:  yes
        Current Load Balance:  round-robin
        Logical Unit Group ID:  NA
        Auto Failback:  on
        Auto Probing:  NA

        Paths:
                Initiator Port Name:  100000109bd30070
                Target Port Name:  20b9d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109bd30070
                Target Port Name:  20b8d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109bd3006f
                Target Port Name:  20b3d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109bd3006f
                Target Port Name:  20b4d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no
```

```
        Target Port Groups:
                ID:  1003
                Explicit Failover:  no
                Access State:  active optimized
                Target Ports:
                        Name:  20b9d039ea593393
                        Relative ID:  8

                        Name:  20b4d039ea593393
                        Relative ID:  3

                ID:  1002
                Explicit Failover:  no
                Access State:  active optimized
                Target Ports:
                        Name:  20b8d039ea593393
                        Relative ID:  7

                        Name:  20b3d039ea593393
                        Relative ID:  2
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
root@chatsol-54-195:~# mpathadm show lu
/dev/rdsk/c0t600A0980383044376C3F4E694E506E44d0s2
Logical Unit:  /dev/rdsk/c0t600A0980383044376C3F4E694E506E44d0s2
        mpath-support:  libmpscsi_vhci.so
        Vendor:  NETAPP
        Product:  LUN C-Mode
        Revision:  9171
        Name Type:  unknown type
        Name:  600a0980383044376c3f4e694e506e44
        Asymmetric:  yes
        Current Load Balance:  round-robin
        Logical Unit Group ID:  NA
        Auto Failback:  on
        Auto Probing:  NA

        Paths:

                Initiator Port Name:  100000109b56c5fb
                Target Port Name:  205200a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109b56c5fb
                Target Port Name:  205000a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
                Path State:  OK
                Demoted:  yes
                Disabled:  no

                Initiator Port Name:  100000109b56c5fa
                Target Port Name:  204f00a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
                Path State:  OK
                Demoted:  yes
                Disabled:  no

                Initiator Port Name:  100000109b56c5fa
                Target Port Name:  205100a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
```

```
                Path State:  OK
                Disabled:  no

        Target Port Groups:
                ID:  1001
                Explicit Failover:  no
                Access State:  active not optimized
                Target Ports:
                        Name:  205200a098ba7afe
                        Relative ID:  8

                        Name:  205100a098ba7afe
                        Relative ID:  7

                ID:  1000
                Explicit Failover:  no
                Access State:  active optimized
                Target Ports:
                        Name:  205000a098ba7afe
                        Relative ID:  6

                        Name:  204f00a098ba7afe
                        Relative ID:  5
```

Step 4: Review the known issues

The Solaris 11.4 release for FCP and iSCSI with ONTAP storage has the following known issues:

NetApp Bug ID	Title	Description	Oracle ID
1362435	HUK 6.2 and Solaris_11.4 FC driver binding changes	Refer to Solaris 11.4 and HUK recommendations. FC driver binding is changed from `ssd (4D)` to `sd (4D)`. Move the existing configuration from `ssd.conf` to `sd.conf` as mentioned in Oracle DOC: 2595926.1). The behavior varies across newly installed Solaris 11.4 systems and systems upgraded from Solaris 11.3 or earlier versions.	(Doc ID 2595926.1)

NetApp Bug ID	Title	Description	Oracle ID
1366780	Solaris LIF issue noticed during storage failover (SFO) giveback operation with Emulex 32G host bus adapter (HBA) on x86 Arch	Solaris LIF issue noticed with Emulex firmware version 12.6.x and later on the x86_64 platform.	SR 3-24746803021
1368957	Solaris 11.x `cfgadm -c configure` resulting in I/O error with end-to-end Emulex configuration	Running `cfgadm -c configure` on Emulex end-to-end configuration results in an I/O error. This is fixed in ONTAP 9.5P17, 9.6P14 , 9.7P13, and 9.8P2	Not Applicable
1345622	Abnormal path reporting on Solaris hosts with ASA/PPorts using OS native commands	Intermittent path reporting issues are noticed on Solaris 11.4 with All SAN Array (ASA).	Not Applicable

What's next?

Learn about using the Solaris Host Utilities tool.

Configure Solaris 11.3 for FCP and iSCSI with ONTAP storage

The Solaris Host Utilities software provides management and diagnostic tools for Solaris hosts that are connected to ONTAP storage. When you install the Solaris Host Utilities on a Solaris 11.3 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability. If your configuration doesn't support SAN booting, you can use a local boot.

SAN boot

SAN booting is the process of setting up a SAN-attached disk (a LUN) as a boot device for a Solaris host. You can set up a SAN boot LUN to work in a Solaris MPxIO environment that is using the FC protocol and running Solaris Host Utilities. The method you use to set up a SAN boot LUN depends on your volume manager and file system.

Steps

1. Use the Interoperability Matrix Tool to verify that your Solaris OS, protocol, and ONTAP version support SAN booting.
2. Follow the best practices for setting up a SAN boot in the Solaris vendor documentation.

Local boot

Perform a local boot by installing the Solaris OS on the local hard disk, for example, install on an SSD, SATA, or RAID.

Step 2: Install the Solaris Host Utilities

NetApp strongly recommends installing the Solaris Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Installing the Solaris Host Utilities changes some of the timeout settings on your Solaris host.

Install Solaris Host Utilities 6.2.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Solaris 11.3 to manage ONTAP LUNs.

Multipathing enables you to configure multiple network paths between the host and storage systems. If one path fails, traffic continues on the remaining paths.

Steps

1. The Solaris Host Utilities load the NetApp recommended parameter settings for SPARC and x86_64 processors.

Show parameter settings

Parameter	Value
throttle_max	8
not_ready_retries	300
busy_retries	30
reset_retries	30
throttle_min	2
timeout_retries	10
physical_block_size	4096
disksort	false
cache-nonvolatile	true

2. If your storage configuration includes MetroCluster, Oracle Solaris virtualization, or SnapMirror active sync, change the default settings:

MetroCluster

By default, the Solaris OS fails to execute the I/O operations after **20s** if all paths to a LUN are lost. This is controlled by the `fcp_offline_delay` parameter. The default value for `fcp_offline_delay` is the appropriate for standard ONTAP clusters. However, in MetroCluster configurations, you must increase the value of `fcp_offline_delay` to **120s** to ensure that I/O doesn't prematurely time out during operations, including unplanned failovers.

For additional information and recommended changes to default settings for MetroCluster, see the Knowledge Base article Solaris host support considerations in a MetroCluster configuration.

Oracle Solaris virtualization

- Solaris virtualization options include Solaris Logical Domains (also called LDOMs or Oracle VM Server for SPARC), Solaris Dynamic Domains, Solaris Zones, and Solaris Containers. These technologies are also referred to as "Oracle Virtual Machines".
- You can use multiple options together, for example, a Solaris Container within a particular Solaris Logical Domain.
- NetApp supports the use of Solaris virtualization technologies where the overall configuration is supported by Oracle and any partition with direct access to LUNs is listed on the IMT in a supported configuration. This includes root containers, LDOM I/O domains, and LDOM using NPIV to access LUNs.
- Partitions or virtual machines that use only virtualized storage resources, such as a `vdsk`, don't need specific qualifications because they don't have direct access to ONTAP LUNs. You only need to verify that the partition or virtual machine that has direct access to the underlying LUN, such as an LDOM I/O domain, is listed on the IMT.

Steps

When LUNs are used as virtual disk devices within an LDOM, the source of the LUN is masked by virtualization and the LDOM doesn't correctly detect the block sizes. To prevent this issue:

a. Patch the LDOM OS for *Oracle Bug 15824910*

b. Create a `vdc.conf` file that sets the block size of the virtual disk to `4096`. See Oracle DOC: 2157669.1 for more information.

c. Verify the patch installation to ensure that the recommended settings have been configured correctly:

   i. Create a zpool:

   ```
   zpool create zpool_name disk_list
   ```

   ii. Run `zdb -C` against the zpool and verify that the value of **ashift** is `12`.

   If the value of **ashift** isn't `12`, re run `zdb -C11`, and verify that the correct patch was installed and recheck the contents of `vdc.conf`.

   Don't proceed until **ashift** shows a value of `12`.

Patches are available for Oracle bug 15824910 on several Solaris versions. Contact Oracle if you need assistance in determining the best kernel patch.

SnapMirror active sync

Beginning with ONTAP 9.9.1, SnapMirror active sync setting configurations are supported in the Solaris host. To verify that the Solaris client applications are non-disruptive when an unplanned site failover switchover occurs in a SnapMirror active sync environment, you must configure the `scsi-vhci-failover-override` setting on the Solaris host. This setting overrides the failover module `f_tpgs` to prevent the execution of the code path that detects the contradiction.

Steps

a. Create the configuration file `/etc/driver/drv/scsi_vhci.conf` with an entry similar to the following example for the NetApp storage type connected to the host:

```
scsi-vhci-failover-override =
"NETAPP  LUN","f_tpgs"
```

b. Verify that the override parameter has been successfully applied:

```
devprop
```

```
mdb
```

Show examples

```
root@host-A:~# devprop -v -n /scsi_vhci scsi-vhci-failover-
override      scsi-vhci-failover-override=NETAPP  LUN + f_tpgs
root@host-A:~# echo "*scsi_vhci_dip::print -x struct dev_info
devi_child | ::list struct dev_info devi_sibling| ::print
struct dev_info devi_mdi_client| ::print mdi_client_t
ct_vprivate| ::print struct scsi_vhci_lun svl_lun_wwn
svl_fops_name"| mdb -k
```

```
svl_lun_wwn = 0xa002a1c8960 "600a0980383134775431524539787938"
svl_fops_name = 0xa00298d69e0 "conf f_tpgs"
```

After `scsi-vhci-failover-override` has been applied, `conf` is added to `svl_fops_name`. For additional information and recommended changes to default settings, refer to the NetApp Knowledge Base article Solaris Host support recommended settings in SnapMirror active sync configuration.

3. Verify that 4KB aligned I/O with zpools using ONTAP LUNs is supported:

a. Verify that your Solaris host is installed with the latest Support Repository Update (SRU):

```
pkg info entire`
```

b. Verify that the ONTAP LUN has `ostype` as "Solaris", independent of the LUN size:

```
lun show -vserver` <vsersver_name>
```

Show example

```
chat-a800-31-33-35-37::*> lun show -vserver solaris_fcp -path
/vol/sol_195_zpool_vol_9/lun -fields ostype
vserver     path                          ostype
----------- ----------------------------- -------
solaris_fcp /vol/sol_195_zpool_vol_9/lun solaris
```

4. Verify the output for your ONTAP LUNs:

```
sanlun lun show
```

You should see an output similar to the following example for an ASA, AFF, or FAS configuration:

Show example

```
root@sparc-s7-55-148:~# sanlun lun show -pv

                    ONTAP Path: Solaris_148_siteA:/vol/Triage/lun
                           LUN: 0
                      LUN Size: 20g
                   Host Device:
/dev/rdsk/c0t600A098038314B32685D573064776172d0s2
                          Mode: C
              Multipath Provider: Sun Microsystems
                Multipath Policy: Native
```

5. Verify the path status for your ONTAP LUNs:

```
mpathadm show lu <LUN>`
```

The following example outputs show the correct path status for ONTAP LUNs in an ASA, AFF, or FAS configuration. The path priorities are displayed against "Access State" for each LUN in the output.

ASA configurations

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
root@sparc-s7-55-82:~# mpathadm show lu
/dev/rdsk/c0t600A098038313953495D58674777794Bd0s2
Logical Unit:  /dev/rdsk/c0t600A098038313953495D58674777794Bd0s2
        mpath-support:  libmpscsi_vhci.so
        Vendor:  NETAPP
        Product:  LUN C-Mode
        Revision:  9171
        Name Type:  unknown type
        Name:  600a098038313953495d58674777794b
        Asymmetric:  yes
        Current Load Balance:  round-robin
        Logical Unit Group ID:  NA
        Auto Failback:  on
        Auto Probing:  NA

        Paths:
                Initiator Port Name:  100000109bd30070
                Target Port Name:  20b9d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109bd30070
                Target Port Name:  20b8d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109bd3006f
                Target Port Name:  20b3d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109bd3006f
                Target Port Name:  20b4d039ea593393
                Logical Unit Number:  0
                Override Path:  NA
                Path State:  OK
                Disabled:  no
```

```
        Target Port Groups:
                ID:  1003
                Explicit Failover:  no
                Access State:  active optimized
                Target Ports:
                        Name:  20b9d039ea593393
                        Relative ID:  8

                        Name:  20b4d039ea593393
                        Relative ID:  3

                ID:  1002
                Explicit Failover:  no
                Access State:  active optimized
                Target Ports:
                        Name:  20b8d039ea593393
                        Relative ID:  7

                        Name:  20b3d039ea593393
                        Relative ID:  2
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the correct output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
root@chatsol-54-195:~# mpathadm show lu
/dev/rdsk/c0t600A0980383044376C3F4E694E506E44d0s2
Logical Unit:  /dev/rdsk/c0t600A0980383044376C3F4E694E506E44d0s2
        mpath-support:  libmpscsi_vhci.so
        Vendor:  NETAPP
        Product:  LUN C-Mode
        Revision:  9171
        Name Type:  unknown type
        Name:  600a0980383044376c3f4e694e506e44
        Asymmetric:  yes
        Current Load Balance:  round-robin
        Logical Unit Group ID:  NA
        Auto Failback:  on
        Auto Probing:  NA

        Paths:

                Initiator Port Name:  100000109b56c5fb
                Target Port Name:  205200a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
                Path State:  OK
                Disabled:  no

                Initiator Port Name:  100000109b56c5fb
                Target Port Name:  205000a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
                Path State:  OK
                Demoted:  yes
                Disabled:  no

                Initiator Port Name:  100000109b56c5fa
                Target Port Name:  204f00a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
                Path State:  OK
                Demoted:  yes
                Disabled:  no

                Initiator Port Name:  100000109b56c5fa
                Target Port Name:  205100a098ba7afe
                Logical Unit Number:  1
                Override Path:  NA
```

```
                Path State:  OK
                Disabled:  no

        Target Port Groups:
                ID:  1001
                Explicit Failover:  no
                Access State:  active not optimized
                Target Ports:
                        Name:  205200a098ba7afe
                        Relative ID:  8

                        Name:  205100a098ba7afe
                        Relative ID:  7

                ID:  1000
                Explicit Failover:  no
                Access State:  active optimized
                Target Ports:
                        Name:  205000a098ba7afe
                        Relative ID:  6

                        Name:  204f00a098ba7afe
                        Relative ID:  5
```

Step 4: Review the known issues

The Solaris 11.3 release for FCP and iSCSI with ONTAP storage has the following known issues:

NetApp Bug ID	Title	Description	Oracle ID
1366780	Solaris LIF problem during GB with Emulex 32G HBA on x86 Arch	Seen with Emulex Firmware version 12.6.x and later on x86_64 Platform	SR 3-24746803021
1368957	Solaris 11.x 'cfgadm -c configure' resulting in I/O error with End-to-End Emulex configuration	Running `cfgadm -c configure` on Emulex end-to-end configurations results in I/O error. This is fixed in ONTAP 9.5P17, 9.6P14, 9.7P13 and 9.8P2	Not Applicable

What's next?

Learn about using the Solaris Host Utilities tool.

SUSE Linux Enterprise Server

Configure SUSE Linux Enterprise Server 16 for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a SUSE Linux Enterprise Server 16 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.
3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 7.1.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with SUSE Linux Enterprise Server 16 to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exists:

   ```
   ls /etc/multipath.conf
   ```

   If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the host operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the native Linux OS compiled multipath parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 3:0:7:9    sdco 69:192  active ready running
  |- 3:0:8:9    sddi 71:0    active ready running
  |- 14:0:8:9   sdjq 65:320  active ready running
  `- 14:0:7:9   sdiw 8:256   active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 3:0:3:0    sdd  8:48    active ready running
| |- 3:0:4:0    sdx  65:112  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 14:0:2:0   sdfk 130:96  active ready running
  `- 14:0:5:0   sdgz 132:240 active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi-2.1.11-160000.2.2.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about SUSE Linux Virtualization (KVM)

  SUSE Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure SUSE Linux Enterprise Server 15 SPx for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on a SUSE Linux Enterprise Server 15 SPx host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with SUSE Linux Enterprise Server 15 SPx to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exists:

   ```
   ls /etc/multipath.conf
   ```

   If the file doesn't exist, create an empty, zero-byte file:

   ```
   touch /etc/multipath.conf
   ```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

   ```
   systemctl enable multipathd
   ```

   ```
   systemctl start multipathd
   ```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the host operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the native Linux OS compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 3:0:7:9    sdco 69:192  active ready running
  |- 3:0:8:9    sddi 71:0    active ready running
  |- 14:0:8:9   sdjq 65:320  active ready running
  `- 14:0:7:9   sdiw 8:256   active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a09803831347657244e527766394e dm-5 NETAPP,LUN C-Mode
size=80G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 3:0:3:0    sdd  8:48    active ready running
| |- 3:0:4:0    sdx  65:112  active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 14:0:2:0   sdfk 130:96  active ready running
  `- 14:0:5:0   sdgz 132:240 active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (iscsi-initiator-utils) is installed:

   ```
   rpm -qa | grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi-2.1.11-160000.2.2.x86_64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
● iscsid.service - Open-iSCSI
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Tue 2025-12-02 11:36:21 EST; 2
weeks 1 day ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
             man:iscsiuio(8)
             man:iscsiadm(8)
   Main PID: 2263 (iscsid)
     Status: "Ready to process requests"
      Tasks: 1 (limit: 816061)
     Memory: 18.5M
        CPU: 14.480s
     CGroup: /system.slice/iscsid.service
             └─2263 /usr/sbin/iscsid -f -d2
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.30.87
192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 -p
192.168.30.87:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.30.87:3260,1139 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [2] 192.168.31.97:3260,1142 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [3] 192.168.31.87:3260,1141 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
tcp: [4] 192.168.30.97:3260,1140 iqn.1992-
08.com.netapp:sn.064a9b19b3ee11f09dcad039eabac370:vs.23 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> ⚠ You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

SUSE Linux Enterprise Server 15 SPx with ONTAP storage has the following known issues.

15 SP1

NetApp Bug ID	Title	Description
1246622	Remote ports transit to a blocked state on SLES15SP1 with Emulex LPe12002 8GB FC during storage failover operations.	Remote ports transit to a blocked state on SLES15SP1 with Emulex LPe12002 8GB Fibre Channel (FC) during storage failover operations. When the storage node returns to an optimal state, the LIFs also come up and the remote port state should read "online." Occasionally, the remote port state might continue to read as "blocked" or "not present." This state can lead to a "failed faulty" path to LUNs at the multipath layer as well as an I/O outage for those LUNs. You can check the remoteport's details against the following sample commands: ---- cat/sys/class/fc_host/host*/device/rport*/fc_remote_ports/rport*/port_name cat/sys/class/fc_host/host*/device/rport*/fc_remote_ports/rport*/port_state ----

15

NetApp Bug ID	Title	Description
1154309	SLES 15 host with more than 20 mapped LUNs might go into maintenance mode after a reboot	SLES 15 host with more than 20 mapped LUNs might go into maintenance mode after a reboot. The maintenance mode becomes single user mode following the message: `Give root password for maintenance (or press Control-D to continue)`

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about SUSE Linux Virtualization (KVM)

  SUSE Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single

physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Ubuntu

Configure Ubuntu 24.04 for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on an Ubuntu 24.04 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Ubuntu 24.04 to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

   Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a0980383145595533f524d6c652f62 dm-24 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:1:13 sdm  8:192  active ready running
  |- 11:0:3:13 sdah 66:16  active ready running
  |- 12:0:1:13 sdbc 67:96  active ready running
  `- 12:0:3:13 sdbx 68:176 active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383148373524536946b542f4a dm-0 NETAPP,LUN C-Mode
size=160G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 14:0:3:0 sdbk 67:224 active ready running
| `- 15:0:2:0 sdbl 67:240 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 14:0:0:0 sda  8:0    active ready running
  `- 15:0:1:0 sdv  65:80  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

   ```
   $apt list |grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
days ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
   Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
             ├─1127418 /usr/sbin/iscsid
             └─1127419 /usr/sbin/iscsid
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Ubuntu Linux Virtualization (KVM)

  Ubuntu Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure Ubuntu 22.04 for FCP and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. When you install the Linux Host Utilities on an Ubuntu 22.04 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Install Linux Host Utilities 8.0.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the multipath configuration for your host

You can use multipathing with Ubuntu 22.04 to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a0980383145595533f524d6c652f62 dm-24 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:1:13 sdm  8:192  active ready running
  |- 11:0:3:13 sdah 66:16  active ready running
  |- 12:0:1:13 sdbc 67:96  active ready running
  `- 12:0:3:13 sdbx 68:176 active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383148373524536946b542f4a dm-0 NETAPP,LUN C-Mode
size=160G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 14:0:3:0 sdbk 67:224 active ready running
| `- 15:0:2:0 sdbl 67:240 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 14:0:0:0 sda  8:0    active ready running
  `- 15:0:1:0 sdv  65:80  active ready running
```

Step 4: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

   ```
   $apt list |grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
days ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
   Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
             ├─1127418 /usr/sbin/iscsid
             └─1127419 /usr/sbin/iscsid
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

Step 5: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 6: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 7: Review the known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.
- Learn about ASM mirroring

  Automatic Storage Management (ASM) mirroring might require changes to the Linux multipath settings to allow ASM to recognize a problem and switch over to an alternate failure group. Most ASM configurations on ONTAP use external redundancy, which means that data protection is provided by the external array and ASM doesn't mirror data. Some sites use ASM with normal redundancy to provide two-way mirroring, normally across different sites. See Oracle databases on ONTAP for further information.

- Learn about Ubuntu Linux Virtualization (KVM)

  Ubuntu Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Configure Ubuntu 20.04 for FCP and iSCSI with ONTAP storage

Configure Ubuntu 20.04 for multipathing and with specific parameters and settings for FCP and iSCSI protocol operations with ONTAP storage.

> The Linux Host Utilities software package doesn't support Ubuntu operating systems.
>
> You don't need to manually configure Kernel-based Virtual Machine (KVM) settings because ONTAP LUNs are automatically mapped to the hypervisor.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Confirm the multipath configuration for your host

You can use multipathing with Ubuntu 20.04 to manage ONTAP LUNs.

To ensure that multipathing is configured correctly for your host, verify that the `/etc/multipath.conf` file is defined and that you have the NetApp recommended settings configured for your ONTAP LUNs.

Steps

1. Verify that the `/etc/multipath.conf` file exits. If the file doesn't exist, create an empty, zero-byte file:

```
touch /etc/multipath.conf
```

2. The first time the `multipath.conf` file is created, you might need to enable and start the multipath services to load the recommended settings:

```
systemctl enable multipathd
```

```
systemctl start multipathd
```

3. Each time you boot the host, the empty `/etc/multipath.conf` zero-byte file automatically loads the NetApp recommended host multipath parameters as the default settings. You shouldn't need to make changes to the `/etc/multipath.conf` file for your host because the operating system is compiled with the multipath parameters that recognize and manage ONTAP LUNs correctly.

   The following table shows the Linux OS native compiled multipath parameter settings for ONTAP LUNs.

Show parameter settings

Parameter	Setting
detect_prio	yes
dev_loss_tmo	"infinity"
failback	immediate
fast_io_fail_tmo	5
features	"2 pg_init_retries 50"
flush_on_last_del	"yes"
hardware_handler	"0"
no_path_retry	queue
path_checker	"tur"
path_grouping_policy	"group_by_prio"
path_selector	"service-time 0"
polling_interval	5
prio	"ontap"
product	LUN
retain_attached_hw_handler	yes
rr_weight	"uniform"
user_friendly_names	no
vendor	NETAPP

4. Verify the parameter settings and path status for your ONTAP LUNs:

```
multipath -ll
```

The default multipath parameters support ASA, AFF, and FAS configurations. In these configurations, a single ONTAP LUN shouldn’t require more than four paths. Having more than four paths can cause problems during a storage failure.

The following example outputs show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# multipath -ll
3600a0980383145595533f524d6c652f62 dm-24 NETAPP,LUN C-Mode
size=10G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
  |- 11:0:1:13 sdm  8:192  active ready running
  |- 11:0:3:13 sdah 66:16  active ready running
  |- 12:0:1:13 sdbc 67:96  active ready running
  `- 12:0:3:13 sdbx 68:176 active ready running
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# multipath -ll
3600a0980383148373524536946b542f4a dm-0 NETAPP,LUN C-Mode
size=160G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=active
| |- 14:0:3:0 sdbk 67:224 active ready running
| `- 15:0:2:0 sdbl 67:240 active ready running
`-+- policy='service-time 0' prio=10 status=enabled
  |- 14:0:0:0 sda  8:0    active ready running
  `- 15:0:1:0 sdv  65:80  active ready running
```

Step 3: Confirm the iSCSI configuration for your host

Ensure that iSCSI is configured correctly for your host.

About this task

You perform the following steps on the iSCSI host.

Steps

1. Verify that the iSCSI initiator package (open-iscsi) is installed:

   ```
   $apt list |grep open-iscsi
   ```

   You should see an output similar to the following example:

   ```
   open-iscsi/noble-updates,noble-updates,now 2.1.9-3ubuntu5.4 amd64
   ```

2. Verify the iSCSI initiator node name, which is located in the `/etc/iscsi/initiatorname.iscsi` file:

   ```
   InitiatorName=iqn.YYYY-MM.com.<vendor>:<host_name>
   ```

3. Configure the iSCSI session timeout parameter located in the `/etc/iscsi/iscsid.conf` file:

   ```
   node.session.timeo.replacement_timeout = 5
   ```

   The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. You should set the value of `replacement_timeout` to 5 in the iSCSI configuration file.

4. Enable the iSCSI service:

   ```
   $systemctl enable iscsid
   ```

5. Start the iSCSI service:

   ```
   $systemctl start iscsid
   ```

6. Verify that the iSCSI service is running:

   ```
   $systemctl status iscsid
   ```

Show example

```
●iscsid.service - iSCSI initiator daemon (iscsid)
     Loaded: loaded (/usr/lib/systemd/system/iscsid.service;
enabled; preset: disabled)
     Active: active (running) since Mon 2026-01-12 12:53:18 IST; 2
days ago
TriggeredBy: ● iscsid.socket
       Docs: man:iscsid(8)
   Main PID: 1127419 (iscsid)
      Tasks: 2 (limit: 76557)
     Memory: 4.3M (peak: 8.8M)
        CPU: 1.657s
     CGroup: /system.slice/iscsid.service
             ├─1127418 /usr/sbin/iscsid
             └─1127419 /usr/sbin/iscsid
```

7. Discover the iSCSI targets:

```
$iscsiadm --mode discovery --op update --type sendtargets --portal
<target_IP>
```

show example

```
iscsiadm --mode discovery --op update --type sendtargets --portal
192.168.100.197
192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8
```

8. Log in to the targets:

```
$iscsiadm --mode node -l all
```

9. Set iSCSI to log in automatically when the host boots:

```
$iscsiadm --mode node -T <target_name> -p <ip:port> -o update -n
node.startup -v automatic
```

You should see an output similar to the following example:

```
iscsiadm --mode node -T iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 -p
192.168.100.197:3260 -o update -n node.startup -v automatic
```

10. Verify the iSCSI sessions:

```
$iscsiadm --mode session
```

Show example

```
iscsiadm --mode session
tcp: [1] 192.168.200.197:3260,1047 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [2] 192.168.100.197:3260,1046 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [3] 192.168.100.199:3260,1048 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
tcp: [4] 192.168.200.199:3260,1049 iqn.1992-
08.com.netapp:sn.7cd154a7d35411f0a25ed039eaa95f59:vs.8 (non-flash)
```

Step 4: Optionally, exclude a device from multipathing

If required, you can exclude a device from multipathing by adding the WWID for the unwanted device to the "blacklist" stanza for the `multipath.conf` file.

Steps

1. Determine the WWID:

```
/lib/udev/scsi_id -gud /dev/sda
```

"sda" is the local SCSI disk that you want to add to the blacklist.

An example WWID is `360030057024d0730239134810c0cb833`.

2. Add the WWID to the "blacklist" stanza:

```
blacklist {
         wwid   360030057024d0730239134810c0cb833
        devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
        devnode "^hd[a-z]"
        devnode "^cciss.*"
}
```

Step 5: Customize multipath parameters for ONTAP LUNs

If your host is connected to LUNs from other vendors and any of the multipath parameter settings are overridden, you need to correct them by adding stanzas later in the `multipath.conf` file that apply specifically to ONTAP LUNs. If you don't do this, the ONTAP LUNs might not work as expected.

Check your `/etc/multipath.conf` file, especially in the defaults section, for settings that might be overriding the default settings for multipath parameters.

> ⚠ You shouldn't override the recommended parameter settings for ONTAP LUNs. These settings are required for optimal performance of your host configuration. Contact NetApp support, your OS vendor, or both for more information.

The following example shows how to correct an overridden default. In this example, the `multipath.conf` file defines values for `path_checker` and `no_path_retry` that aren't compatible with ONTAP LUNs, and you can't remove these parameters because ONTAP storage arrays are still attached to the host. Instead, you correct the values for `path_checker` and `no_path_retry` by adding a device stanza to the `multipath.conf` file that applies specifically to the ONTAP LUNs.

Show example

```
defaults {
   path_checker       readsector0
   no_path_retry      fail
}

devices {
   device {
      vendor           "NETAPP"
      product          "LUN"
      no_path_retry    queue
      path_checker     tur
   }
}
```

Step 6: Review the known issues

There are no known issues.

What's next?

- Learn about Ubuntu Linux Virtualization (KVM)

  Ubuntu Linux can serve as a KVM host. This enables you to run multiple virtual machines on a single physical server using the Linux Kernel-based Virtual Machine (KVM) technology. The KVM host doesn't require explicit host configuration settings for ONTAP LUNs.

Veritas

Configure Veritas Infoscale 9 for FC, FCoE, and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. Use the Linux Host Utilities with Veritas Infoscale 9 for Oracle Linux (Red Hat Compatible Kernel based), Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server hosts to support managing FC, FCoE, and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

- Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.
- Refer to the Veritas Support Portal (Product Matrix, Platform Lookup, and HCL Matrix) to verify the SAN boot configuration support and known issues.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the Veritas Dynamic Multipathing configuration for your host

Use Veritas Dynamic Multipathing (VxDMP) with Veritas Infoscale 9 to manage ONTAP LUNs.

To ensure that VxDMP is set up correctly for your host, you need to verify the VxDMP configuration and check the Array Support Library (ASL) and the Array Policy Module (APM) configuration. The ASL and APM

packages for NetApp storage systems are installed during the Veritas software installation.

> For heterogenous multipath environments, including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, refer to the Veritas Product Administration documentation for the configuration settings.

Before you begin

Ensure your configuration meets the system requirements. See the Interoperability Matrix Tool and the Veritas HCL Matrix.

Steps

1. Verify the ONTAP target array is attached to the VxDMP multipath:

```
vxdmpadm
```

Show example

```
#vxdmpadm listenclosure
ENCLR_NAME        ENCLR_TYPE     ENCLR_SNO      STATUS
ARRAY_TYPE     LUN_COUNT    FIRMWARE
================================================================================
================================
info_asa0         Info_ASA       81KDT+YTg35P         CONNECTED
ALUA       20          9161
infoscal1         Infoscal       81Ocq?Z7hPzC         CONNECTED
ALUA       23          9181
# vxdmpadm getdmpnode
NAME       STATE    ENCLR-TYPE  PATHS  ENBL  DSBL ENCLR-NAME
================================================================================
infoscal1_22         ENABLED      Infoscal     4      4     0
infoscal1
```

2. Check the configuration of the ASL and APM packages. NetApp recommends that you use the latest supported packages listed on the Veritas support portal.

Show example ASL and APM configuration

```
#  vxdmpadm list dmpnode dmpnodename=infoscal1_22 | grep asl
asl             = libvxnetapp.so

# vxddladm listversion |grep libvxnetapp.so
libvxnetapp.so                  vm-8.0.0-rev-1   8.0

# rpm -qa |grep VRTSaslapm
VRTSaslapm-9.0.3-RHEL9.x86_64
#vxddladm listsupport libname=libvxnetapp.so
ATTR_NAME   ATTR_VALUE
=========================================
LIBNAME     libvxnetapp.so
VID         NETAPP
PID         All
ARRAY_TYPE  ALUA, A/A
```

3. For optimum system configuration in storage failover operations, verify that you have the following Veritas VxDMP tunables:

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

4. Set the DMP tunables to online:

```
# vxdmpadm settune dmp_tunable=value
```

5. Verify that the tunables settings are correct:

```
# vxdmpadm gettune
```

The following example shows the effective VxDMP tunables on a SAN host.

Show example

```
# vxdmpadm gettune

Tunable                    Current Value    Default Value
dmp_cache_open                        on               on
dmp_daemon_count                      10               10
dmp_delayq_interval                   15               15
dmp_display_alua_states               on               on
dmp_fast_recovery                     on               on
dmp_health_time                       60               60
dmp_iostats_state                enabled          enabled
dmp_log_level                          1                1
dmp_low_impact_probe                  on               on
dmp_lun_retry_timeout                 60               30
dmp_path_age                         120              300
dmp_pathswitch_blks_shift              9                9
dmp_probe_idle_lun                    on               on
dmp_probe_threshold                    5                5
dmp_restore_cycles                    10               10
dmp_restore_interval                  60              300
dmp_restore_policy         check_disabled   check_disabled
dmp_restore_state                enabled          enabled
dmp_retry_count                        5                5
dmp_scsi_timeout                      20               20
dmp_sfg_threshold                      1                1
dmp_stat_interval                      1                1
dmp_monitor_ownership                 on               on
dmp_monitor_fabric                    on               on
dmp_native_support                   off              off
```

6. Configure the protocol timeout values:

FC/FCoE

Use the default timeout values for FC and FCoE.

iSCSI

Set the `replacement_timeout` parameter value to 120.

The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. NetApp recommends setting the value of `replacement_timeout` to 120 in the iSCSI configuration file.

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

7. Verify the parameter settings and path status for your ONTAP LUNs:

   In AFF, FAS, or ASA configurations, a single ONTAP LUN should not require more than four paths. More than four paths can cause problems during a storage failure.

   The following examples show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# vxdmpadm getsubpaths dmpnodename=infoscal1_21
NAME   STATE[A]    PATH-TYPE[M] CTLR-NAME  ENCLR-TYPE  ENCLR-
NAME   ATTRS  PRIORITY
=============================================================
==================
sdby   ENABLED(A)  Active/Optimized  c1    Infoscal    infoscal1
-      -
sddx   ENABLED(A)  Active/Optimized  c2    Infoscal    infoscal1
-      -
sdfe   ENABLED(A)  Active/Optimized  c1    Infoscal    infoscal1
-      -
sdfo   ENABLED(A)  Active/Optimized  c2    Infoscal    infoscal1
-      -
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME   STATE[A]    PATH-TYPE[M]    CTLR-NAME    ENCLR-TYPE   ENCLR-
NAME   ATTRS   PRIORITY
=============================================================
==================
sdas   ENABLED     Active/Non-Optimized c13    SFRAC        sfrac0
-       -
sdb    ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdcj   ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdea   ENABLED     Active/Non-Optimized c14    SFRAC        sfrac0
-       -
```

Step 4: Known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.

Configure Veritas Infoscale 8 for FC, FCoE, and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. Use the Linux Host Utilities with Veritas Infoscale 8 for Oracle Linux (Red Hat Compatible Kernel based), Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server hosts to support managing FC, FCoE, and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

- Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.
- Refer to the Veritas Support Portal (Product Matrix, Platform Lookup, and HCL Matrix) to verify the SAN boot configuration support and known issues.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

> Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the Veritas Dynamic Multipathing configuration for your host

Use Veritas Dynamic Multipathing (VxDMP) with Veritas Infoscale 8 to manage ONTAP LUNs.

To ensure that VxDMP is set up correctly for your host, you need to verify the VxDMP configuration and check the Array Support Library (ASL) and the Array Policy Module (APM) configuration. The ASL and APM packages for NetApp storage systems are installed during the Veritas software installation.

> For heterogenous multipath environments, including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, refer to the Veritas Product Administration documentation for the configuration settings.

Before you begin

Ensure that your configuration meets the system requirements. See the Interoperability Matrix Tool and the Veritas HCL Matrix.

Steps

1. Verify the ONTAP target array is attached to the VxDMP multipath:

```
vxdmpadm
```

Show example

```
# vxdmpadm listenclosure
ENCLR_NAME    ENCLR_TYPE  ENCLR_SNO      STATUS       ARRAY_TYPE
LUN_COUNT     FIRMWARE
============================================================================
=================
sfrac0        SFRAC       804Xw$PqE52h   CONNECTED    ALUA           43
9800
# vxdmpadm getdmpnode
NAME        STATE     ENCLR-TYPE   PATHS   ENBL   DSBL ENCLR-NAME
============================================================================
sfrac0_47   ENABLED   SFRAC        4       4      0    sfrac0
```

2. Check the configuration of the ASL and APM packages. NetApp recommends that you use the latest supported packages listed on the Veritas support portal.

   Show example ASL and APM configuration

```
# vxdmpadm list dmpnode dmpnodename=sfrac0_47 | grep asl
asl        = libvxnetapp.so
# vxddladm listversion |grep libvxnetapp.so
libvxnetapp.so                  vm-8.0.0-rev-1   8.0

# rpm -qa |grep VRTSaslapm
VRTSaslapm-x.x.x.0000-RHEL8.X86_64
vxddladm listsupport libname=libvxnetapp.so
ATTR_NAME   ATTR_VALUE
=======================================
LIBNAME     libvxnetapp.so
VID         NETAPP
PID         All
ARRAY_TYPE  ALUA, A/A
```

3. For optimum system configuration in storage failover operations, verify that you have the following Veritas VxDMP tunables:

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

4. Set the DMP tunables to online:

```
# vxdmpadm settune dmp_tunable=value
```

5. Verify that the tunables settings are correct:

```
# vxdmpadm gettune
```

   The following example shows the effective VxDMP tunables on a SAN host.

Show example

```
# vxdmpadm gettune

Tunable                     Current Value    Default Value
dmp_cache_open                         on               on
dmp_daemon_count                       10               10
dmp_delayq_interval                    15               15
dmp_display_alua_states                on               on
dmp_fast_recovery                      on               on
dmp_health_time                        60               60
dmp_iostats_state                 enabled          enabled
dmp_log_level                           1                1
dmp_low_impact_probe                   on               on
dmp_lun_retry_timeout                  60               30
dmp_path_age                          120              300
dmp_pathswitch_blks_shift               9                9
dmp_probe_idle_lun                     on               on
dmp_probe_threshold                     5                5
dmp_restore_cycles                     10               10
dmp_restore_interval                   60              300
dmp_restore_policy          check_disabled   check_disabled
dmp_restore_state                 enabled          enabled
dmp_retry_count                         5                5
dmp_scsi_timeout                       20               20
dmp_sfg_threshold                       1                1
dmp_stat_interval                       1                1
dmp_monitor_ownership                  on               on
dmp_monitor_fabric                     on               on
dmp_native_support                    off              off
```

6. Configure the protocol timeout values:

FC/FCoE

Use the default timeout values for FC and FCoE.

iSCSI

Set the `replacement_timeout` parameter value to 120.

The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. NetApp recommends setting the value of `replacement_timeout` to 120 in the iSCSI configuration file.

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

7. Set the host "udev rport" values for RHEL 8 and 9 series hosts to support the Veritas Infoscale environment in storage failover scenarios.

   Configure "udev rport" values by creating the file `/etc/udev/rules.d/40-rport.rules` with the following file content:

```
# cat /etc/udev/rules.d/40-rport.rules
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add",
RUN+=/bin/sh -c 'echo 20 >
/sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000
>/sys/class/fc_remote_ports/%k/dev_loss_tmo'"
```

   Refer to the standard Veritas Infoscale product documentation for all other settings specific to Veritas.

8. Verify the parameter settings and path status for your ONTAP LUNs:

   In AFF, FAS, or ASA configurations, a single ONTAP LUN should not require more than four paths. More than four paths can cause problems during a storage failure.

   The following examples show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME  STATE[A]   PATH-TYPE[M]   CTLR-NAME   ENCLR-TYPE  ENCLR-
NAME  ATTRS  PRIORITY
=============================================================
==================
sdas  ENABLED (A)    Active/Optimized c13   SFRAC       sfrac0
-      -
sdb   ENABLED(A) Active/Optimized     c14   SFRAC       sfrac0
-      -
sdcj  ENABLED(A)  Active/Optimized     c14   SFRAC       sfrac0
-      -
sdea  ENABLED (A)    Active/Optimized c14   SFRAC       sfrac0
-
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME   STATE[A]    PATH-TYPE[M]    CTLR-NAME    ENCLR-TYPE   ENCLR-
NAME   ATTRS   PRIORITY
=============================================================
==================
sdas   ENABLED     Active/Non-Optimized c13    SFRAC        sfrac0
-       -
sdb    ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdcj   ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdea   ENABLED     Active/Non-Optimized c14    SFRAC        sfrac0
-       -
```

Step 4: Known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.

Configure Veritas Infoscale 7 for FC, FCoE, and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. Use the Linux Host Utilities with Veritas Infoscale 7 for Oracle Linux (Red Hat Compatible Kernel based), Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server hosts to support managing FC, FCoE, and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

- Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.
- Refer to the Veritas Support Portal (Product Matrix, Platform Lookup, and HCL Matrix) to verify the SAN boot configuration support and known issues.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the Veritas Dynamic Multipathing configuration for your host

Use Veritas Dynamic Multipathing (VxDMP) with Veritas Infoscale 7 to manage ONTAP LUNs.

To ensure that VxDMP is set up correctly for your host, you need to verify the VxDMP configuration and check the Array Support Library (ASL) and the Array Policy Module (APM) configuration. The ASL and APM packages for NetApp storage systems are installed during the Veritas software installation.

For heterogenous multipath environments, including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, refer to the Veritas Product Administration documentation for the configuration settings.

Before you begin

Ensure that your configuration meets the system requirements. See the Interoperability Matrix Tool and the Veritas HCL Matrix.

Steps

1. Verify the ONTAP target array is attached to the VxDMP multipath:

```
vxdmpadm
```

Show example

```
# vxdmpadm listenclosure
ENCLR_NAME    ENCLR_TYPE  ENCLR_SNO      STATUS       ARRAY_TYPE
LUN_COUNT     FIRMWARE
=============================================================================
=================
sfrac0        SFRAC       804Xw$PqE52h   CONNECTED    ALUA           43
9800
# vxdmpadm getdmpnode
NAME        STATE     ENCLR-TYPE   PATHS   ENBL   DSBL ENCLR-NAME
=============================================================================
sfrac0_47   ENABLED   SFRAC        4       4      0    sfrac0
```

2. Check the configuration of the ASL and APM packages. NetApp recommends that you use the latest supported packages listed on the Veritas support portal.

   Show example ASL and APM configuration

```
# vxdmpadm list dmpnode dmpnodename=sfrac0_47 | grep asl
asl        = libvxnetapp.so
# vxddladm listversion |grep libvxnetapp.so
libvxnetapp.so                  vm-8.0.0-rev-1   8.0

# rpm -qa |grep VRTSaslapm
VRTSaslapm-x.x.x.0000-RHEL8.X86_64
vxddladm listsupport libname=libvxnetapp.so
ATTR_NAME   ATTR_VALUE
=======================================
LIBNAME     libvxnetapp.so
VID         NETAPP
PID         All
ARRAY_TYPE  ALUA, A/A
```

3. For optimum system configuration in storage failover operations, verify that you have the following Veritas VxDMP tunables:

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

4. Set the DMP tunables to online:

```
# vxdmpadm settune dmp_tunable=value
```

5. Verify that the tunables settings are correct:

```
# vxdmpadm gettune
```

   The following example shows the effective VxDMP tunables on a SAN host.

Show example

```
# vxdmpadm gettune

Tunable                     Current Value    Default Value
dmp_cache_open                         on               on
dmp_daemon_count                       10               10
dmp_delayq_interval                    15               15
dmp_display_alua_states                on               on
dmp_fast_recovery                      on               on
dmp_health_time                        60               60
dmp_iostats_state                 enabled          enabled
dmp_log_level                           1                1
dmp_low_impact_probe                   on               on
dmp_lun_retry_timeout                  60               30
dmp_path_age                          120              300
dmp_pathswitch_blks_shift               9                9
dmp_probe_idle_lun                     on               on
dmp_probe_threshold                     5                5
dmp_restore_cycles                     10               10
dmp_restore_interval                   60              300
dmp_restore_policy          check_disabled   check_disabled
dmp_restore_state                 enabled          enabled
dmp_retry_count                         5                5
dmp_scsi_timeout                       20               20
dmp_sfg_threshold                       1                1
dmp_stat_interval                       1                1
dmp_monitor_ownership                  on               on
dmp_monitor_fabric                     on               on
dmp_native_support                    off              off
```

6. Configure the protocol timeout values:

FC/FCoE

Use the default timeout values for FC and FCoE.

iSCSI

Set the `replacement_timeout` parameter value to 120.

The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. NetApp recommends setting the value of `replacement_timeout` to 120 in the iSCSI configuration file.

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

7. Set the host "udev rport" values for RHEL 8 and 9 series hosts to support the Veritas Infoscale environment in storage failover scenarios.

   Configure "udev rport" values by creating the file `/etc/udev/rules.d/40-rport.rules` with the following file content:

```
# cat /etc/udev/rules.d/40-rport.rules
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add",
RUN+=/bin/sh -c 'echo 20 >
/sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000
>/sys/class/fc_remote_ports/%k/dev_loss_tmo'"
```

   Refer to the standard Veritas Infoscale product documentation for all other settings specific to Veritas.

8. Verify the parameter settings and path status for your ONTAP LUNs:

   In AFF, FAS, or ASA configurations, a single ONTAP LUN should not require more than four paths. More than four paths can cause problems during a storage failure.

   The following examples show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME  STATE[A]   PATH-TYPE[M]   CTLR-NAME   ENCLR-TYPE  ENCLR-
NAME  ATTRS  PRIORITY
=============================================================
==================
sdas  ENABLED (A)    Active/Optimized c13   SFRAC       sfrac0
-      -
sdb   ENABLED(A) Active/Optimized     c14   SFRAC       sfrac0
-      -
sdcj  ENABLED(A)  Active/Optimized     c14   SFRAC       sfrac0
-      -
sdea  ENABLED (A)    Active/Optimized c14   SFRAC       sfrac0
-
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME   STATE[A]    PATH-TYPE[M]    CTLR-NAME    ENCLR-TYPE   ENCLR-
NAME   ATTRS   PRIORITY
=============================================================
==================
sdas   ENABLED     Active/Non-Optimized c13    SFRAC        sfrac0
-       -
sdb    ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdcj   ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdea   ENABLED     Active/Non-Optimized c14    SFRAC        sfrac0
-       -
```

Step 4: Known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.

Configure Veritas Infoscale 6 for FC, FCoE, and iSCSI with ONTAP storage

The Linux Host Utilities software provides management and diagnostic tools for Linux hosts that are connected to ONTAP storage. Use the Linux Host Utilities with Veritas Infoscale 6 for Oracle Linux (Red Hat Compatible Kernel based), Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server hosts to support managing FC, FCoE, and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can configure your host to use SAN booting to simplify deployment and improve scalability.

Before you begin

- Use the Interoperability Matrix Tool to verify that your Linux OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.
- Refer to the Veritas Support Portal (Product Matrix, Platform Lookup, and HCL Matrix) to verify the SAN boot configuration support and known issues.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Step 2: Install the Linux Host Utilities

NetApp strongly recommends installing the Linux Host Utilities to support ONTAP LUN management and assist technical support with gathering configuration data.

> Installing the Linux Host Utilities doesn't change any host timeout settings on your Linux host.

Step 3: Confirm the Veritas Dynamic Multipathing configuration for your host

Use Veritas Dynamic Multipathing (VxDMP) with Veritas Infoscale 6 to manage ONTAP LUNs.

To ensure that VxDMP is set up correctly for your host, you need to verify the VxDMP configuration and check the Array Support Library (ASL) and the Array Policy Module (APM) configuration. The ASL and APM packages for NetApp storage systems are installed during the Veritas software installation.

> For heterogenous multipath environments, including Veritas Infoscale, Linux Native Device Mapper, and LVM volume manager, refer to the Veritas Product Administration documentation for the configuration settings.

Before you begin

Ensure that your configuration meets the system requirements. See the Interoperability Matrix Tool and the Veritas HCL Matrix.

Steps

1. Verify the ONTAP target array is attached to the VxDMP multipath:

```
vxdmpadm
```

Show example

```
# vxdmpadm listenclosure
ENCLR_NAME    ENCLR_TYPE  ENCLR_SNO      STATUS       ARRAY_TYPE
LUN_COUNT     FIRMWARE
============================================================================
=================
sfrac0        SFRAC       804Xw$PqE52h   CONNECTED    ALUA           43
9800
# vxdmpadm getdmpnode
NAME        STATE     ENCLR-TYPE   PATHS   ENBL   DSBL ENCLR-NAME
============================================================================
sfrac0_47   ENABLED   SFRAC        4       4      0    sfrac0
```

2. Check the configuration of the ASL and APM packages. NetApp recommends that you use the latest supported packages listed on the Veritas support portal.

   Show example ASL and APM configuration

```
# vxdmpadm list dmpnode dmpnodename=sfrac0_47 | grep asl
asl        = libvxnetapp.so
# vxddladm listversion |grep libvxnetapp.so
libvxnetapp.so                  vm-8.0.0-rev-1   8.0

# rpm -qa |grep VRTSaslapm
VRTSaslapm-x.x.x.0000-RHEL8.X86_64
vxddladm listsupport libname=libvxnetapp.so
ATTR_NAME   ATTR_VALUE
=======================================
LIBNAME     libvxnetapp.so
VID         NETAPP
PID         All
ARRAY_TYPE  ALUA, A/A
```

3. For optimum system configuration in storage failover operations, verify that you have the following Veritas VxDMP tunables:

Parameter	Setting
dmp_lun_retry_timeout	60
dmp_path_age	120
dmp_restore_interval	60

4. Set the DMP tunables to online:

```
# vxdmpadm settune dmp_tunable=value
```

5. Verify that the tunables settings are correct:

```
# vxdmpadm gettune
```

   The following example shows the effective VxDMP tunables on a SAN host.

Show example

```
# vxdmpadm gettune

Tunable                    Current Value    Default Value
dmp_cache_open                        on               on
dmp_daemon_count                      10               10
dmp_delayq_interval                   15               15
dmp_display_alua_states               on               on
dmp_fast_recovery                     on               on
dmp_health_time                       60               60
dmp_iostats_state                enabled          enabled
dmp_log_level                          1                1
dmp_low_impact_probe                  on               on
dmp_lun_retry_timeout                 60               30
dmp_path_age                         120              300
dmp_pathswitch_blks_shift              9                9
dmp_probe_idle_lun                    on               on
dmp_probe_threshold                    5                5
dmp_restore_cycles                    10               10
dmp_restore_interval                  60              300
dmp_restore_policy         check_disabled   check_disabled
dmp_restore_state                enabled          enabled
dmp_retry_count                        5                5
dmp_scsi_timeout                      20               20
dmp_sfg_threshold                      1                1
dmp_stat_interval                      1                1
dmp_monitor_ownership                 on               on
dmp_monitor_fabric                    on               on
dmp_native_support                   off              off
```

6. Configure the protocol timeout values:

FC/FCoE

Use the default timeout values for FC and FCoE.

iSCSI

Set the `replacement_timeout` parameter value to 120.

The iSCSI `replacement_timeout` parameter controls how long the iSCSI layer should wait for a timed-out path or session to reestablish itself before failing any commands on it. NetApp recommends setting the value of `replacement_timeout` to 120 in the iSCSI configuration file.

```
# grep replacement_timeout /etc/iscsi/iscsid.conf
node.session.timeo.replacement_timeout = 120
```

7. Set the host "udev rport" values for RHEL 8 and 9 series hosts to support the Veritas Infoscale environment in storage failover scenarios.

   Configure "udev rport" values by creating the file `/etc/udev/rules.d/40-rport.rules` with the following file content:

```
# cat /etc/udev/rules.d/40-rport.rules
KERNEL=="rport-*", SUBSYSTEM=="fc_remote_ports", ACTION=="add",
RUN+=/bin/sh -c 'echo 20 >
/sys/class/fc_remote_ports/%k/fast_io_fail_tmo;echo 864000
>/sys/class/fc_remote_ports/%k/dev_loss_tmo'"
```

   Refer to the standard Veritas Infoscale product documentation for all other settings specific to Veritas.

8. Verify the parameter settings and path status for your ONTAP LUNs:

   In AFF, FAS, or ASA configurations, a single ONTAP LUN should not require more than four paths. More than four paths can cause problems during a storage failure.

   The following examples show the correct parameter settings and path status for ONTAP LUNs in an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration optimizes all paths to a given LUN, keeping them active. This improves performance by serving I/O operations through all paths at the same time.

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME  STATE[A]   PATH-TYPE[M]   CTLR-NAME   ENCLR-TYPE  ENCLR-
NAME  ATTRS  PRIORITY
=============================================================
==================
sdas  ENABLED (A)    Active/Optimized c13   SFRAC       sfrac0
-      -
sdb   ENABLED(A) Active/Optimized     c14   SFRAC       sfrac0
-      -
sdcj  ENABLED(A)  Active/Optimized     c14   SFRAC       sfrac0
-      -
sdea  ENABLED (A)    Active/Optimized c14   SFRAC       sfrac0
-
```

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with higher and lower priorities. Higher priority Active/Optimized paths are served by the controller where the aggregate is located. Lower priority paths are active but non-optimized because they are served by a different controller. Non-optimized paths are only used when optimized paths aren’t available.

The following example displays the output for an ONTAP LUN with two Active/Optimized paths and two Active/Non-Optimized paths:

Show example

```
# vxdmpadm getsubpaths dmpnodename-sfrac0_47
NAME   STATE[A]    PATH-TYPE[M]    CTLR-NAME    ENCLR-TYPE   ENCLR-
NAME   ATTRS   PRIORITY
=============================================================
==================
sdas   ENABLED     Active/Non-Optimized c13    SFRAC        sfrac0
-       -
sdb    ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdcj   ENABLED(A)  Active/Optimized     c14    SFRAC        sfrac0
-       -
sdea   ENABLED     Active/Non-Optimized c14    SFRAC        sfrac0
-       -
```

Step 4: Known issues

There are no known issues.

What's next?

- Learn about using the Linux Host Utilities tool.

Windows

Configure Windows Server 2025 for FCP and iSCSI with ONTAP storage

The Windows Host Utilities are a set of software programs with documentation that enables you to connect Windows hosts to virtual disks (LUNs) on a NetApp SAN. When you install the Windows Host Utilities on a Windows Server 2025 host, you can use the Host Utilities to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.
.

Step 1: Optionally, enable SAN booting

You can boot the Windows OS using a local boot or a SAN boot. NetApp recommends using a SAN boot to simplify deployment and improve scalability..

SAN boot

If you choose to use SAN booting, it must be supported by your configuration.

Before you begin

Use the Interoperability Matrix Tool to verify that your Windows OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Local boot

Perform a local boot by installing the Windows OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install Windows hotfixes

NetApp recommends installing the **latest cumulative update** available from the Microsoft Update Catalog on the host server.

Steps

1. Download the hotfixes from the Microsoft Update Catalog 2025.

   You need to contact Microsoft support for the hotfixes that aren't available for download from the Microsoft Update Catalog.

2. Follow the instructions provided by Microsoft to install the hotfixes.

   Many hotfixes require a Windows host reboot. You can wait to reboot the host until *after* you install or upgrade the Host Utilities.

Step 3: Install the Windows Host Utilities

The Windows Host Utilities are a set of software programs with documentation that enable you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installing the latest Windows Host Utilities to support ONTAP LUN management and help technical support collect configuration data.

For Windows Host Utilities configuration and installation information, see the Windows Host Utilities documentation and select the installation procedure for your Windows Host Utilities version.

Step 4: Confirm the multipath configuration for your host

Install the Microsoft Multipath I/O (MPIO) software and enable multipathing if your Windows host has more than one path to the storage system.

On a Windows system, the two main components in an MPIO solution are the device-specific module (DSM) and the Windows MPIO. MPIO presents one disk to the Windows OS for all paths and the DSM manages path failovers.

If you don't install the MPIO software, the Windows OS might see each path as a separate disk. This can lead to data corruption.

Windows XP or Windows Vista running in a Hyper-V virtual machine doesn't support MPIO.

Steps

1. Install the MPIO software and enable multipathing.
2. When you select MPIO on systems using FC, the Host Utilities installer sets the required timeout values for Emulex and QLogic FC HBAs.

Emulex FC

The timeout values for Emulex FC HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

QLogic FC

The timeout values for QLogic FC HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

3. Verify the path status for your ONTAP LUNs:

   Depending on your SAN configuration, the host uses ASA, AFF, or FAS configurations to access ONTAP LUNs. These configurations shouldn't require more than four paths to access a single ONTAP LUN. More than four paths can cause problems during storage failure.

   The following example outputs show the correct settings for ONTAP LUNs for an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration should have one group of Active/Optimized paths with single priorities. The controller services the paths and sends I/O on all active paths.

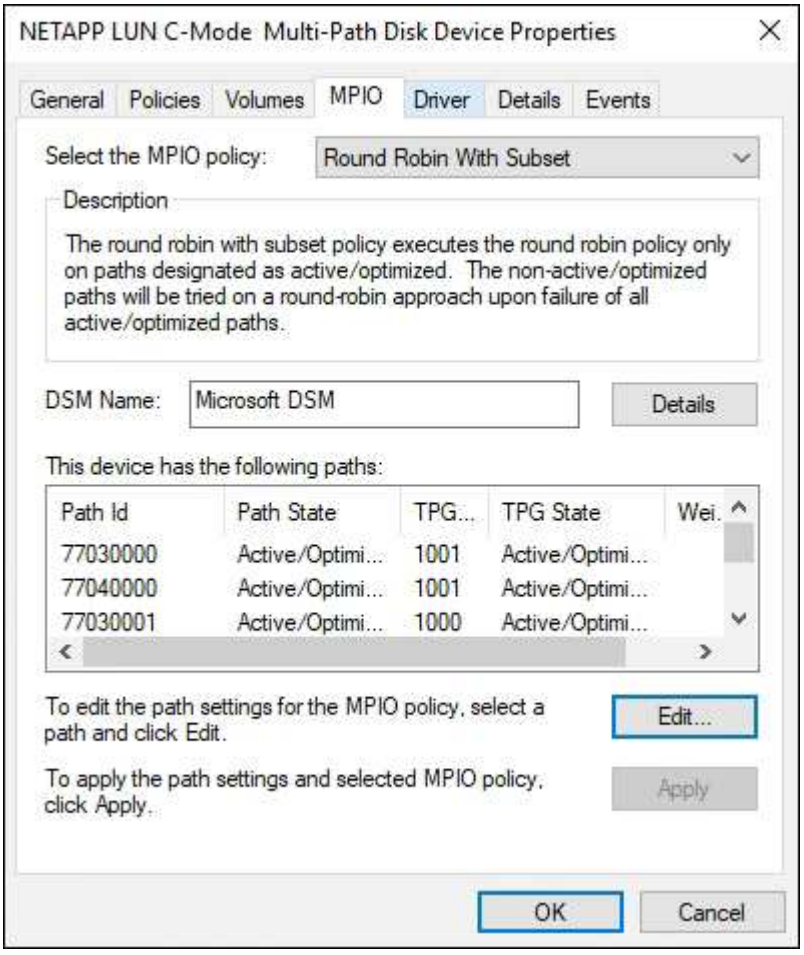

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with different priorities. The paths with higher priorities are Active/Optimized and are serviced by the controller where the aggregate is located. The paths with lower priorities are serviced from a different controller. They are active but non-optimized and are only used when optimized paths aren’t available.

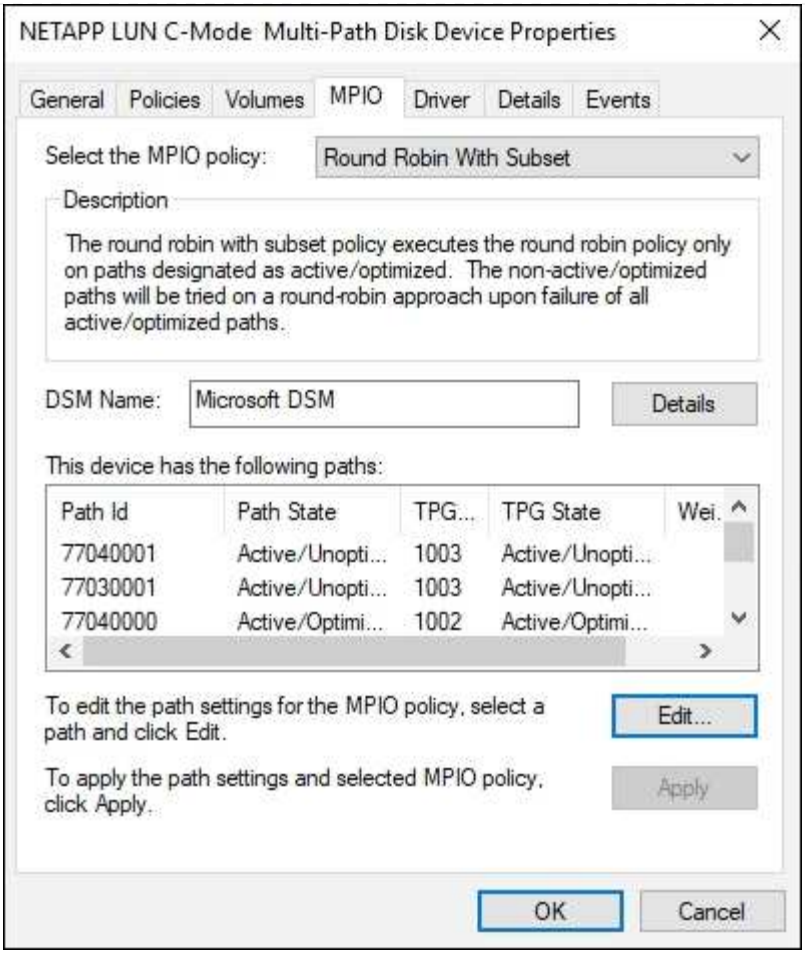

Step 5: Review the Known issues

There are no known issues.

What's next?

Learn about the Windows Host Utilities configuration for ONTAP storage

Configure Windows Server 2022 for FCP and iSCSI with ONTAP storage

The Windows Host Utilities enable you to connect Windows hosts to virtual disks (LUNs) on a NetApp SAN. Install the Windows Host Utilities on a Windows Server 2022 host to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can boot the Windows OS using a local boot or a SAN boot. NetApp recommends using a SAN boot to simplify deployment and improve scalability..

SAN boot

If you choose to use SAN booting, it must be supported by your configuration.

Before you begin

Use the Interoperability Matrix Tool to verify that your Windows OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Local boot

Perform a local boot by installing the Windows OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install Windows hotfixes

NetApp recommends installing the **latest cumulative update** available from the Microsoft Update Catalog on the host server.

Steps

1. Download the hotfixes from the Microsoft Update Catalog 2022.

   You need to contact Microsoft support for the hotfixes that aren't available for download from the Microsoft Update Catalog.

2. Follow the instructions provided by Microsoft to install the hotfixes.

   Many hotfixes require a Windows host reboot. You can wait to reboot the host until *after* you install or upgrade the Host Utilities.

Step 3: Install the Windows Host Utilities

The Windows Host Utilities are a set of software programs with documentation that enable you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installing the latest Windows Host Utilities to support ONTAP LUN management and help technical support collect configuration data.

For Windows Host Utilities configuration and installation information, see the Windows Host Utilities documentation and select the installation procedure for your Windows Host Utilities version.

Step 4: Confirm the multipath configuration for your host

Install the Microsoft Multipath I/O (MPIO) software and enable multipathing if your Windows host has more than one path to the storage system.

On a Windows system, the two main components in an MPIO solution are the device-specific module (DSM) and the Windows MPIO. MPIO presents one disk to the Windows OS for all paths and the DSM manages path failovers.

If you don't install the MPIO software, the Windows OS might see each path as a separate disk. This can lead to data corruption.

Windows XP or Windows Vista running in a Hyper-V virtual machine doesn't support MPIO.

Steps

1. Install the MPIO software and enable multipathing.
2. When you select MPIO on systems using FC, the Host Utilities installer sets the required timeout values for Emulex and QLogic FC HBAs.

Emulex FC

The timeout values for Emulex FC HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

QLogic FC

The timeout values for QLogic FC HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

3. Verify the path status for your ONTAP LUNs:

   Depending on your SAN configuration, the host uses ASA, AFF, or FAS configurations to access ONTAP LUNs. These configurations shouldn't require more than four paths to access a single ONTAP LUN. More than four paths can cause problems during storage failure.

   The following example outputs show the correct settings for ONTAP LUNs for an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration should have one group of Active/Optimized paths with single priorities. The controller services the paths and sends I/O on all active paths.

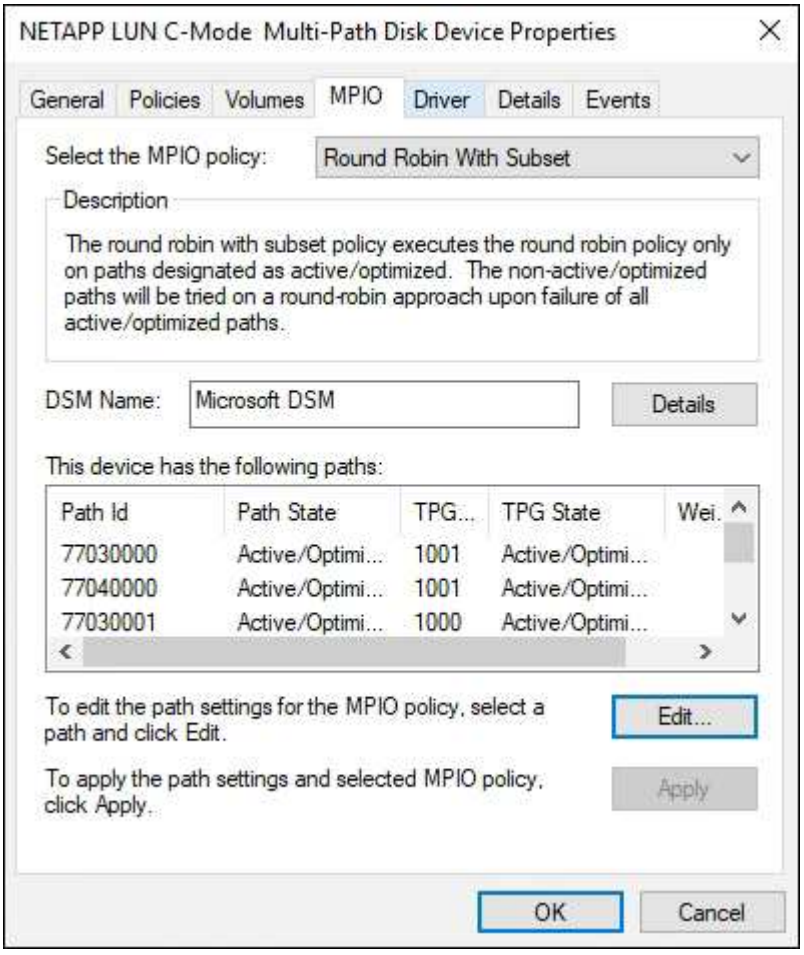

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with different priorities. The paths with higher priorities are Active/Optimized and are serviced by the controller where the aggregate is located. The paths with lower priorities are serviced from a different controller. They are active but non-optimized and are only used when optimized paths aren’t available.

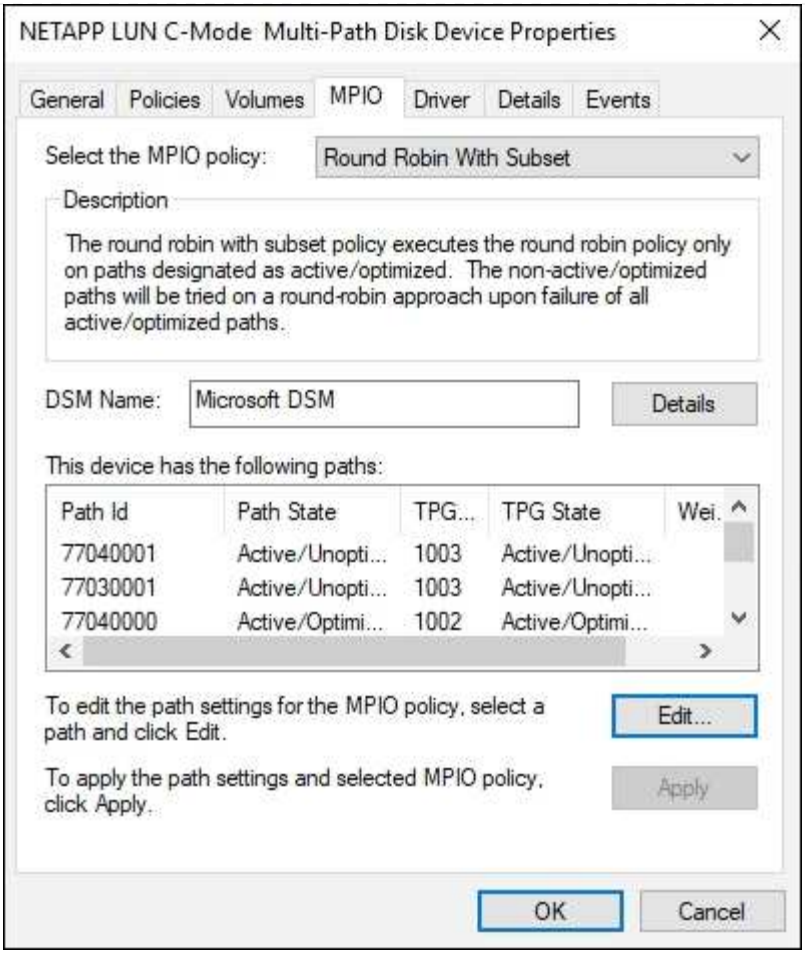

Step 5: Review the Known issues

There are no known issues.

What's next?

Learn about the Windows Host Utilities configuration for ONTAP storage

Configure Windows Server 2019 for FCP and iSCSI with ONTAP storage

The Windows Host Utilities enable you to connect Windows hosts to virtual disks (LUNs) on a NetApp SAN. Install the Windows Host Utilities on a Windows Server 2019 host to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can boot the Windows OS using a local boot or a SAN boot. NetApp recommends using a SAN boot to simplify deployment and improve scalability..

SAN boot

If you choose to use SAN booting, it must be supported by your configuration.

Before you begin

Use the Interoperability Matrix Tool to verify that your Windows OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Local boot

Perform a local boot by installing the Windows OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install Windows hotfixes

NetApp recommends installing the **latest cumulative update** available from the Microsoft Update Catalog on the host server.

Steps

1. Download the hotfixes from the Microsoft Update Catalog 2019.

   You need to contact Microsoft support for the hotfixes that aren't available for download from the Microsoft Update Catalog.

2. Follow the instructions provided by Microsoft to install the hotfixes.

   Many hotfixes require a Windows host reboot. You can wait to reboot the host until *after* you install or upgrade the Host Utilities.

Step 3: Install the Windows Host Utilities

The Windows Host Utilities are a set of software programs with documentation that enable you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installing the latest Windows Host Utilities to support ONTAP LUN management and help technical support collect configuration data.

For Windows Host Utilities configuration and installation information, see the Windows Host Utilities documentation and select the installation procedure for your Windows Host Utilities version.

Step 4: Confirm the multipath configuration for your host

Install the Microsoft Multipath I/O (MPIO) software and enable multipathing if your Windows host has more than one path to the storage system.

On a Windows system, the two main components in an MPIO solution are the device-specific module (DSM) and the Windows MPIO. MPIO presents one disk to the Windows OS for all paths and the DSM manages path failovers.

If you don't install the MPIO software, the Windows OS might see each path as a separate disk. This can lead to data corruption.

Windows XP or Windows Vista running in a Hyper-V virtual machine doesn't support MPIO.

Steps

1. Install the MPIO software and enable multipathing.
2. When you select MPIO on systems using FC, the Host Utilities installer sets the required timeout values for Emulex and QLogic FC HBAs.

Emulex FC

The timeout values for Emulex FC HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

QLogic FC

The timeout values for QLogic FC HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

3. Verify the path status for your ONTAP LUNs:

   Depending on your SAN configuration, the host uses ASA, AFF, or FAS configurations to access ONTAP LUNs. These configurations shouldn't require more than four paths to access a single ONTAP LUN. More than four paths can cause problems during storage failure.

   The following example outputs show the correct settings for ONTAP LUNs for an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration should have one group of Active/Optimized paths with single priorities. The controller services the paths and sends I/O on all active paths.

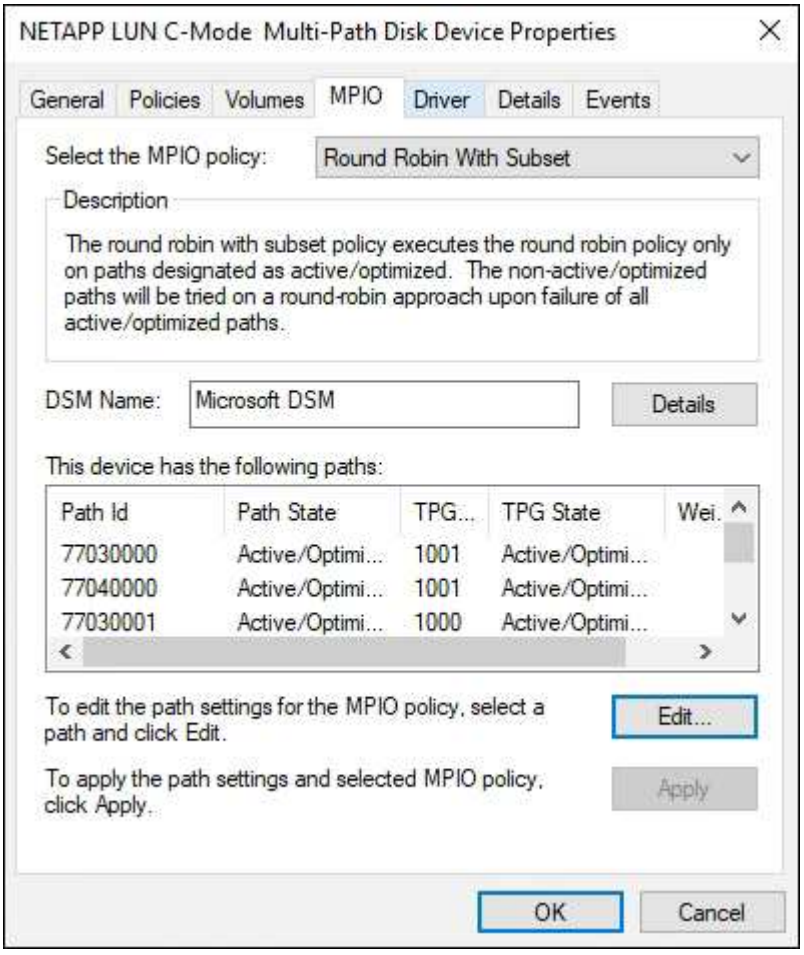

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with different priorities. The paths with higher priorities are Active/Optimized and are serviced by the controller where the aggregate is located. The paths with lower priorities are serviced from a different controller. They are active but non-optimized and are only used when optimized paths aren't available.

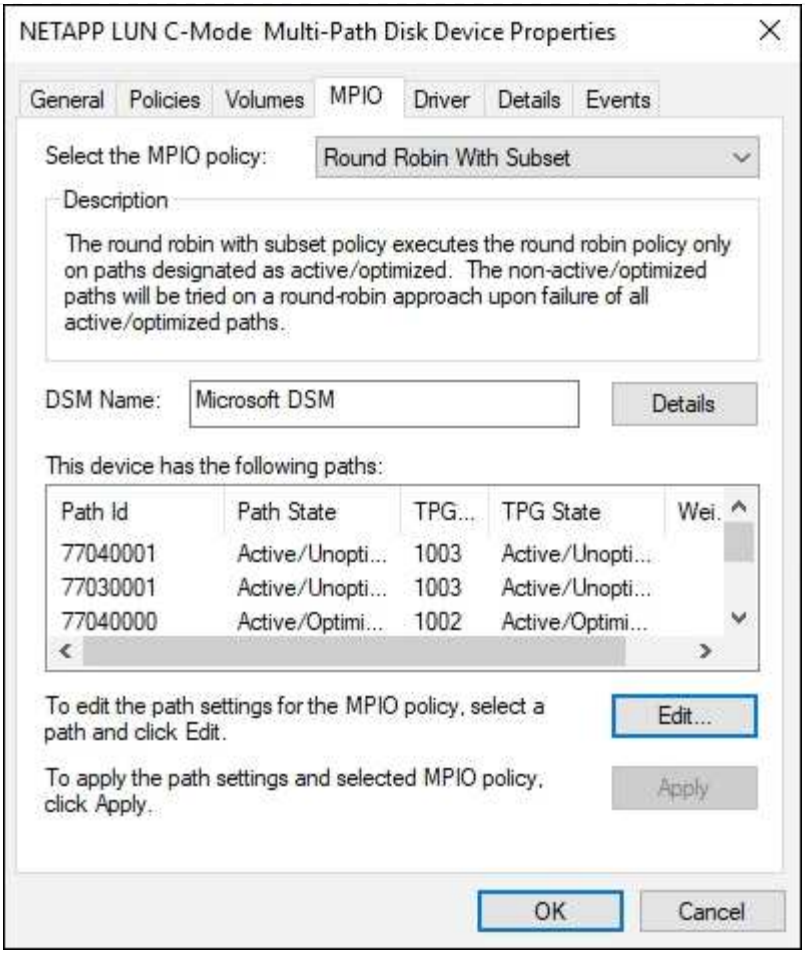

Step 5: Review the Known issues

There are no known issues.

What's next?

Learn about the Windows Host Utilities configuration for ONTAP storage

Configure Windows Server 2016 for FCP and iSCSI with ONTAP storage

The Windows Host Utilities enable you to connect Windows hosts to virtual disks (LUNs) on a NetApp SAN. Install the Windows Host Utilities on a Windows Server 2016 host to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can boot the Windows OS using a local boot or a SAN boot. NetApp recommends using a SAN boot to simplify deployment and improve scalability..

SAN boot

If you choose to use SAN booting, it must be supported by your configuration.

Before you begin

Use the Interoperability Matrix Tool to verify that your Windows OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Local boot

Perform a local boot by installing the Windows OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install Windows hotfixes

NetApp recommends installing the **latest cumulative update** available from the Microsoft Update Catalog on the host server.

Steps

1. Download the hotfixes from the Microsoft Update Catalog 2016.

   You need to contact Microsoft support for the hotfixes that aren't available for download from the Microsoft Update Catalog.

2. Follow the instructions provided by Microsoft to install the hotfixes.

   Many hotfixes require a Windows host reboot. You can wait to reboot the host until *after* you install or upgrade the Host Utilities.

Step 3: Install the Windows Host Utilities

The Windows Host Utilities are a set of software programs with documentation that enable you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installing the latest Windows Host Utilities to support ONTAP LUN management and help technical support collect configuration data.

For Windows Host Utilities configuration and installation information, see the Windows Host Utilities documentation and select the installation procedure for your Windows Host Utilities version.

Step 4: Confirm the multipath configuration for your host

Install the Microsoft Multipath I/O (MPIO) software and enable multipathing if your Windows host has more than one path to the storage system.

On a Windows system, the two main components in an MPIO solution are the device-specific module (DSM) and the Windows MPIO. MPIO presents one disk to the Windows OS for all paths and the DSM manages path failovers.

If you don't install the MPIO software, the Windows OS might see each path as a separate disk. This can lead to data corruption.

Windows XP or Windows Vista running in a Hyper-V virtual machine doesn't support MPIO.

Steps

1. Install the MPIO software and enable multipathing.
2. When you select MPIO on systems using FC, the Host Utilities installer sets the required timeout values for Emulex and QLogic FC HBAs.

Emulex FC

The timeout values for Emulex FC HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

QLogic FC

The timeout values for QLogic FC HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

3. Verify the path status for your ONTAP LUNs:

   Depending on your SAN configuration, the host uses ASA, AFF, or FAS configurations to access ONTAP LUNs. These configurations shouldn't require more than four paths to access a single ONTAP LUN. More than four paths can cause problems during storage failure.

   The following example outputs show the correct settings for ONTAP LUNs for an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration should have one group of Active/Optimized paths with single priorities. The controller services the paths and sends I/O on all active paths.

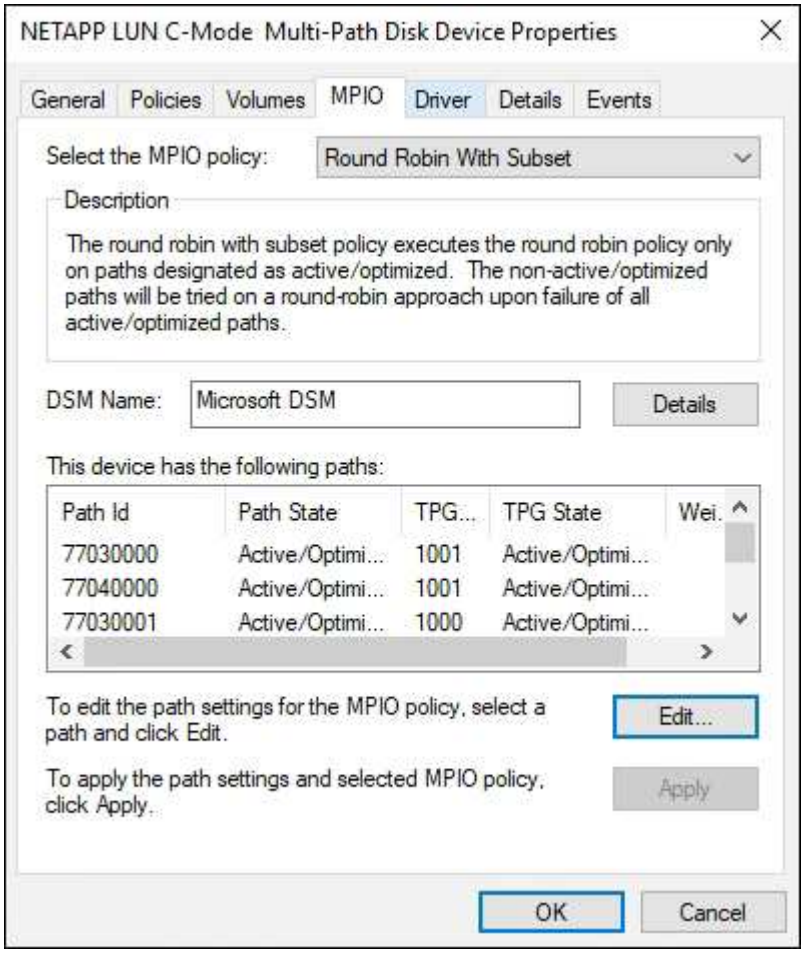

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with different priorities. The paths with higher priorities are Active/Optimized and are serviced by the controller where the aggregate is located. The paths with lower priorities are serviced from a different controller. They are active but non-optimized and are only used when optimized paths aren’t available.

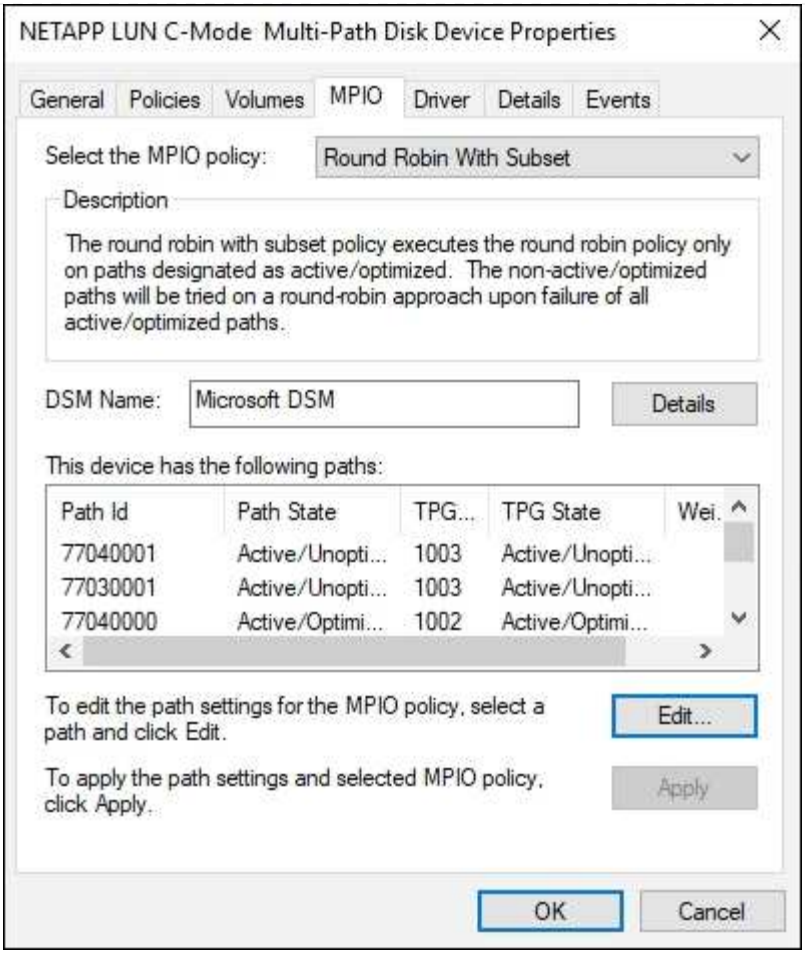

Step 5: Review the Known issues

There are no known issues.

What's next?

Learn about the Windows Host Utilities configuration for ONTAP storage

Configure Windows Server 2012 R2 for FCP and iSCSI with ONTAP storage

The Windows Host Utilities enable you to connect Windows hosts to virtual disks (LUNs) on a NetApp SAN. Install the Windows Host Utilities on a Windows Server 2012 R2 host to help you manage FCP and iSCSI protocol operations with ONTAP LUNs.

Step 1: Optionally, enable SAN booting

You can boot the Windows OS using a local boot or a SAN boot. NetApp recommends using a SAN boot to simplify deployment and improve scalability..

SAN boot

If you choose to use SAN booting, it must be supported by your configuration.

Before you begin

Use the Interoperability Matrix Tool to verify that your Windows OS, host bus adapter (HBA), HBA firmware, HBA boot BIOS, and ONTAP version support SAN booting.

Steps

1. Create a SAN boot LUN and map it to the host.
2. Enable SAN booting in the server BIOS for the ports to which the SAN boot LUN is mapped.

   For information on how to enable the HBA BIOS, see your vendor-specific documentation.

3. Verify that the configuration was successful by rebooting the host and verifying that the OS is up and running.

Local boot

Perform a local boot by installing the Windows OS on the local hard disk, for example, on an SSD, SATA, or RAID.

Step 2: Install Windows hotfixes

NetApp recommends installing the **latest cumulative update** available from the Microsoft Update Catalog on the host server.

Steps

1. Download the hotfixes from the Microsoft Update Catalog 2012 R2.

   You need to contact Microsoft support for the hotfixes that aren't available for download from the Microsoft Update Catalog.

2. Follow the instructions provided by Microsoft to install the hotfixes.

   Many hotfixes require a Windows host reboot. You can wait to reboot the host until *after* you install or upgrade the Host Utilities.

Step 3: Install the Windows Host Utilities

The Windows Host Utilities are a set of software programs with documentation that enable you to connect host computers to virtual disks (LUNs) on a NetApp SAN. NetApp recommends downloading and installing the latest Windows Host Utilities to support ONTAP LUN management and help technical support collect configuration data.

For Windows Host Utilities configuration and installation information, see the Windows Host Utilities documentation and select the installation procedure for your Windows Host Utilities version.

Step 4: Confirm the multipath configuration for your host

Install the Microsoft Multipath I/O (MPIO) software and enable multipathing if your Windows host has more than one path to the storage system.

On a Windows system, the two main components in an MPIO solution are the device-specific module (DSM) and the Windows MPIO. MPIO presents one disk to the Windows OS for all paths and the DSM manages path failovers.

If you don't install the MPIO software, the Windows OS might see each path as a separate disk. This can lead to data corruption.

Windows XP or Windows Vista running in a Hyper-V virtual machine doesn't support MPIO.

Steps

1. Install the MPIO software and enable multipathing.
2. When you select MPIO on systems using FC, the Host Utilities installer sets the required timeout values for Emulex and QLogic FC HBAs.

Emulex FC

The timeout values for Emulex FC HBAs:

Property type	Property value
LinkTimeOut	1
NodeTimeOut	10

QLogic FC

The timeout values for QLogic FC HBAs:

Property type	Property value
LinkDownTimeOut	1
PortDownRetryCount	10

3. Verify the path status for your ONTAP LUNs:

   Depending on your SAN configuration, the host uses ASA, AFF, or FAS configurations to access ONTAP LUNs. These configurations shouldn't require more than four paths to access a single ONTAP LUN. More than four paths can cause problems during storage failure.

   The following example outputs show the correct settings for ONTAP LUNs for an ASA, AFF, or FAS configuration.

ASA configuration

An ASA configuration should have one group of Active/Optimized paths with single priorities. The controller services the paths and sends I/O on all active paths.

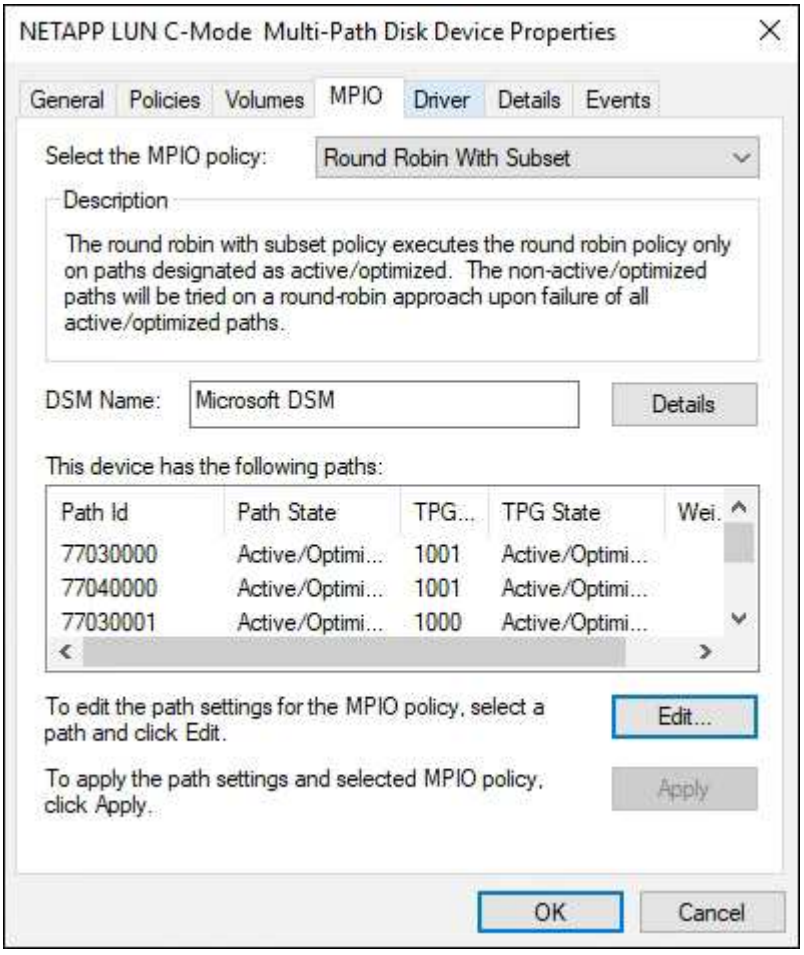

AFF or FAS configuration

An AFF or FAS configuration should have two groups of paths with different priorities. The paths with higher priorities are Active/Optimized and are serviced by the controller where the aggregate is located. The paths with lower priorities are serviced from a different controller. They are active but non-optimized and are only used when optimized paths aren’t available.

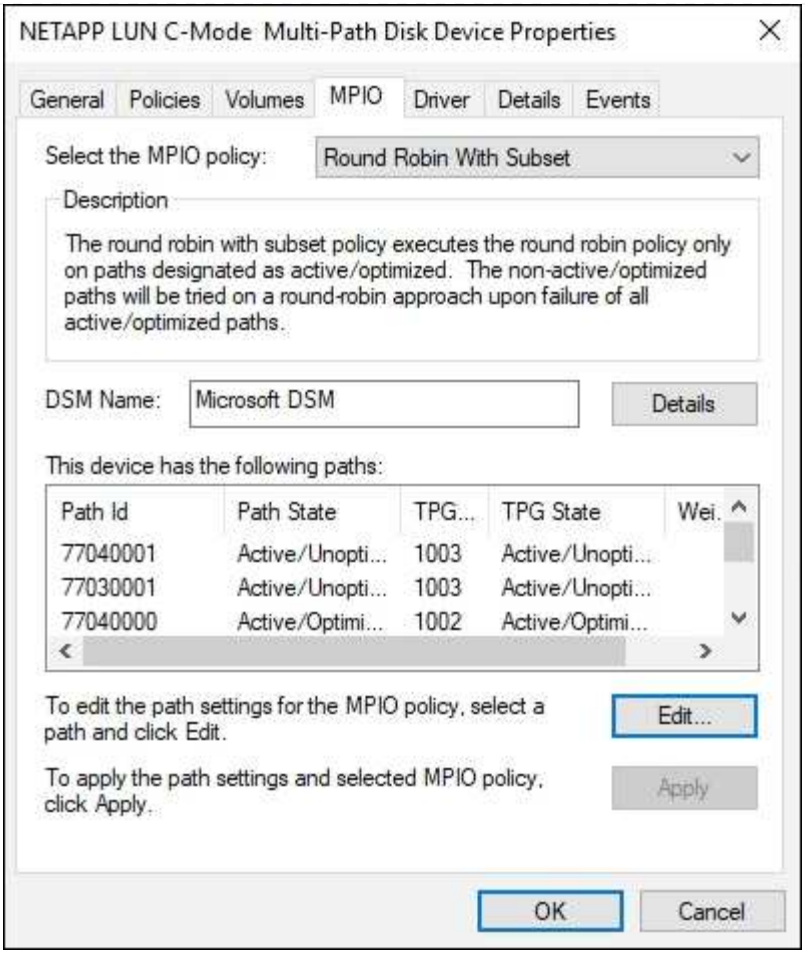

Step 5: Review the Known issues

There are no known issues.

What's next?

Learn about the Windows Host Utilities configuration for ONTAP storage

Copyright information

Trademark information